Abstract of "An Algorithmic Theory of General Equilibrium", by Denizalp Göktaş, Ph.D., Brown University, May 2025.

In the wake of the Second World War, as they presided over major public expenditure reforms, European and American governments supported the development of rigorous mathematical models of economies to guide economic policy. Over the next two decades, general equilibrium models (or Walrasian economies) emerged as the dominant framework. However, as these models were often analytically intractable, as early as the 1960s, a group of researchers led by Herbert Scarf turned their attention to finding "a general method for the explicit numerical solution of [general equilibrium models]." While some methods had limited success in solving simple models, 50 years later, a general method for computing solutions to more complex models remains elusive. Nevertheless, these models—and their often inaccurate solution methods—continue to be widely used in applications such as resource allocation and public policy analysis, raising concerns about the impact of inaccurate solutions on the public good. This thesis addresses this issue by leveraging tools from computer science and game theory to analyze algorithms for general equilibrium models.

The first part of this thesis focuses on variational inequalities (VIs), a mathematical modeling paradigm, and their application to Walrasian economies (i.e., models driven by demand and supply). The second part of this thesis concerns pseudo-games, a multiagent optimization framework, and their application to Arrow-Debreu economies (i.e., Walrasian economies in which demand and supply are generated by consumers and firms). The final part of this thesis concerns Markov pseudo-games, a multiagent *stochastic* optimization framework, and their application to Radner economies (i.e., a generalization of Walrasian and Arrow-Debreu economies that explicitly model time and uncertainty). While Parts 1 and 2 of this thesis resolve Scarf's challenge by solving general equilibrium models developed during his lifetime, Part 3 merely scratches the surface of a new research direction. Above all, it raises an exciting and contemporary analog to Scarf's challenge at the intersection of deep learning, reinforcement learning, and mathematical economics—namely, finding a general method for the explicit numerical solution of Radner economies.

BROWN UNIVERSITY

COMPUTER SCIENCE DEPARTMENT

An Algorithmic Theory of General Equilibrium

DENIZALP GÖKTAŞ

B.A. in Economics-Political Science, Paris Institute of Political Studies, 2019

B.A. in Statistics-Computer Science, Columbia University, 2019

M.S. in Computer Science, Brown University, 2023

SUPERVISOR:

Amy Greenwald, Brown University

THESIS COMMITTEE:

Amy Greenwald, Brown University (Chair)

Simina Brânzei, Purdue University

Richard Cole, New York University

Yu Cheng, Brown University

Geoffrey Gordon, Carnegie Mellon University

Christian Kroer, Columbia University

DISSERTATION SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY.

MAY 2025

KEY WORDS: ARTIFICIAL INTELLIGENCE, COMPUTATIONAL COMPLEXITY, DEEP LEARNING, GAME THEORY, GENERAL EQUILIBRIUM THEORY, MULTIAGENT LEARNING, OPTIMIZATION, VARIATIONAL INEQUALITIES

This dissertation by Denizalp Göktaş is accepted in its present form
by the Department of Computer Science as satisfying the
dissertation requirement for the degree of Doctor of Philosophy.

Date 4/30/2025 ____________________

Signed by: Amy Greenwald
9D67CB6215154CE...

Amy Greenwald, Advisor

Recommended to the Graduate Council

Date 4/30/2025 ____________________

DocuSigned by: Simina Brânzei
10915001292F47C...

Simina Brânzei, Reader

Date 4/29/2025 ____________________

Signed by: Richard Cole
6D0670525B83485...

Richard Cole, Reader

Date 4/29/2025 ____________________

DocuSigned by: Yu Cheng
318EF1F356744C5...

Yu Cheng, Reader

Date 4/29/2025 ____________________

Signed by: Geoffrey Gordon
2F531E74468F4BD...

Geoffrey Gordon, Reader

Date 4/29/2025 ____________________

Signed by: Christian Kroer
911CA896D478497...

Christian Kroer, Reader

Approved by the Graduate Council

Date ____________________ ____________________

Thomas A. Lewis, Dean of the Graduate School

Curriculum Vitae

Denizalp Göktaş holds dual Bachelor of Arts degrees in Computer Science & Statistics from Columbia University, and in Political Science & Economics from the Paris Institute of Political Studies (Sciences Po). In 2019, he began his doctoral studies in Computer Science at Brown University, where he also earned a Master of Science en route to his Ph.D. During his Ph.D., he was a Visiting Scholar at the Simons Institute for the Theory of Computing, a Research Scientist Intern with Google DeepMind's Game Theory and Multiagent team, and a Research Scientist Intern at JPMorgan Chase & Co.'s Multiagent Learning and Simulation Research Lab. He was awarded the JP Morgan AI Fellowship in 2022. Denizalp co-founded the Workshop on Foundation Models, LLMs, and Game Theory (2023) and the Workshop on Computational Methods for Economic Dynamics (2024).

During his Ph.D., he has published the following papers: *BoLT: Building on Local Trust to Solve Lending Market Failure* (Goldstein et al., 2020); *A Consumer-Theoretic Characterization of Fisher Market Equilibria* (Goktas et al., 2022a); *Convex-Concave Min-Max Stackelberg Games* (Goktas and Greenwald, 2021); *Gradient Descent Ascent in Min-Max Stackelberg Games* (Goktas and Greenwald, 2022c); *Robust No-Regret Learning in Min-Max Stackelberg Games* (Goktas and Greenwald, 2022a); *An Algorithmic Theory of Markets and their Application to Decentralized Markets* (Goktas, 2022); *Zero-Sum Stochastic Stackelberg Games* (Goktas et al., 2022c); *Exploitability Minimization in Games and Beyond* (Goktas and Greenwald, 2022b); *Fisher Markets with Social Influence* (Zhao et al., 2023); *Tâtonnement in Homothetic Fisher Markets* (Goktas et al., 2023c); *Convex-Concave Zero-Sum Stochastic Stackelberg Games* (Goktas et al., 2023b); *Generative Adversarial Equilibrium Solvers* (Goktas et al., 2023a); *Banzhaf Power in Hierarchical Games* (Randolph et al., 2024); *STA-RLHF: Stackelberg Aligned Reinforcement Learning with Human Feedback Authors* (Makar-Limanov et al., 2024); *Efficient Inverse Multiagent Learning* (Goktas et al., 2024).

I pursued this thesis as one pursues a work of art.

Dedicated to my mother Deniz
whose love turns the heartlines in my hands into golden threads of fate.
Thank you for everything.

Acknowledgments

One thing my PhD never lacked was adventure. My journey started in the midst of a global pandemic and took me from San Francisco to Berlin, Vienna, London, and Paris, giving me the privilege of spending time with brilliant researchers at the Simons Institute, Google DeepMind, and JP Morgan.

The architect of these opportunities was, without a doubt, my advisor Amy Greenwald. Amy, the woman you are! Compliments would not do you justice, so I'll spare you them. During my first year, the day after a dramatic encounter in the CIT elevator, you said, "Perhaps it's not a coincidence that they call the PhD defense a *viva* in the UK—your task is to defend your research with your life!" I carried those words with me for the rest of my PhD. The only word I can find to describe the last six years together is *wild*, and I am sure we would not have it any other way. Thank you for listening to me and for caring about what I had to say; it is the greatest gift one can receive. I will forever carry the memories and lessons of our almost religious, hours-long conversations on game theory; our long walks through New York, Berkeley, London, New Haven, and Providence; you grilling me in our reading group meetings; and the exacting paper submission deadlines we faced together. *At the start of my PhD, I had one idea of what the advisor-advisee relationship is and should be, now at the end of it, I see that I got that and everything more. Thank you, thank you for everything.*

The inspiration behind my first research projects on general equilibrium were the works of Richard Cole, Yun Kuen (Marco) Cheung, Simina Brânzei, and Christian Kroer. Few people remain in the world with the attention span to catch a mathematical error in Equation 132 of a 50-page paper; Richard and Marco are among them. Thank you for being the trailblazers, giving me endless feedback, and inspiring me to be a better researcher. I wrote every single word in this thesis thinking, "what would Richard think"? Simina, thanks for the numerous chats at Berkeley, hearing your stories, and seeing your passion for research take you all around the world is an inspiration. Christian, thank you for spending time on my committee, we did not get a chance to work closely, but knowing how highly your students speak of you, it speaks volumes to how we should have. At this point, I should also acknowledge two other committee members who have generously joined the committee for this thesis, Yu Cheng and Geoffrey Gordon. Thank you for your time, efforts, and interest in my growth, I am grateful.

I would not have applied to a PhD program if it weren't for two people: Seth Goldstein and Ryan Zohar. Seth, in many ways, this thesis was written to answer the very questions we sought to explore in that first-ever research project we started together. Thank you for inspiring me—without you, I would never have had the chance to write it, and for that, I am deeply grateful. Ryan, my curious friend, your passion for learning, your desire for discovery, and your encouragement to apply for the PhD are why I pursued it in the first place. I cannot wait to see where life takes us.

One of the most exciting opportunities in my life was working at DeepMind in London right after the pandemic. I owe this opportunity to two people: Andrea Tacchetti and Kalesha Bullard. Andrea, thank you for giving me the chance to work with you—I am forever indebted. Your kindness, patience, brilliance, and resilience in expanding into new research horizons, whether you realize it or not, have taught me invaluable lessons that I carry with me. *Thank you*. Kalesha, thank you for preparing me for my interviews, for the time spent together in London's restaurants and markets, and for sharing your stories with me. You are a treasure. You must know that Minae and I refer to you as *the icon*, and I know you know you are one.

Once I arrived at DeepMind, an incredibly talented group of researchers became my champions: David Parkes, Guy Lever, Ian Gemp, Georgios Piliouras, Luke Marris, Yoram Bachrach, Romuald Élie, Minae Kwon, Robert Lange, Evgenii Nikishin, and Chris Lu. David, thank you for supporting my growth and taking an interest in my research. But most importantly, thank you for being a resilient role model—you should know it means the world to me. Guy, my rock, your sense of humor, your perspective on the world, and your steadfast sense of identity made me feel at home at DeepMind. Thank you. Ian, of all the researchers I have met, you are the one with the purest and most genuine love for research. Learning from you and sharing ideas with you was an immense and inspiring experience. Thank you. Georgios, your research creativity, and desire to explore topics at the edge of the field has inspired much of this thesis starting from Berkeley all the way to London, thank you for sharing your ideas and inspirations with me. Luke, you taught me so much about machine learning engineering—thank you for your tireless help. Yoram and Romuald, thank you both for the many hours spent chatting with me, pointing me toward exciting research directions, and supporting me at DeepMind. Minae, you brilliant, noble, and extraordinary woman, I am so grateful to know you and call you my friend. Robert, you loving man, in another life, I know we were together. Evgenii, dreamer, keep sailing toward new horizons—you will only continue to blossom. Chris, kind soul, I am excited to see the great things you will build.

Once back in Providence from London, two talented and resilient researchers from Caltech reached out to me with a question about a paper not included in this thesis: Apurva Badithela and Josefine Graebener. This led to a two-year collaboration on a fascinating project using game theory to solve automated test design. Apurva and Josefine, thank you for sharing your ideas with me and for teaching me about your field.

Before leaving Providence for London to join DeepMind, I was awarded a JP Morgan AI fellowship and put in touch with Sumitra Ganesh to be her mentee. Back from DeepMind, I was hoping to do research at a place with a focus on economics. So, I reached out to Sumitra to ask if I could join her team of researchers for the summer, and she kindly agreed. Sumitra, thank you for giving me research freedom and supporting my research, I am grateful. At JP Morgan, my greatest champion was my research mentor Alec Koppel. Alec, you and the world, knows you are a brilliant and prolific researcher, but to me what mattered was your love for life, your support for those you mentor, and your unbeatable small-talk game. The paper we wrote at JP Morgan was not included in this already far-too-long thesis, but it was

a paper I had dreamed of writing for a very long time, and I believe it is one of the most important papers I published in my PhD. *Thank you for trusting my vision, and supporting this research with no hesitation even in the face of setbacks, you are a role-model.* At JP Morgan, I worked with other brilliant researcher: Sujay Bhatt, Alan Mishler, Kevin Stanghl, Thomas Cook.

Once I left JP Morgan, I set my sights on accomplishing the dream I entered grad school with: building economic simulations with general equilibrium theory. With the goal to find the right literature, I learned more about macroeconomics and stumbled on the works of two great economists: Jesse Perla and Adrien Auclert. Jesse, thank you for caring, being open to listen to and read my ideas. It meant a lot to me at a moment where I felt all alone trying to make sense of the literature. Adrien, thank you for giving me the opportunity to participate in your NBER workshop on Heterogeneous-Agent Macroeconomics, I know non-economists rarely get the chance to do so. Learning about your research and communities was pivotal in concluding the final part of this thesis. I am hopeful that in the near future, computer scientists and economists, will join their forces to solve pressing macroeconomic problems.

In the midst of the never-ending cycle of paper submissions, I was fortunate to meet a brilliant and unique student who has become my greatest supporter and research partner, Sadie Zhao. Sadie, I cannot believe it has been four years since we met! Seeing you grow from an undergraduate interested in research to a PhD student with a research agenda, and now a friend, has been such a joy. *Your future is bright, dream on, and reach for the stars!* The last part of this thesis was a collaboration with Sadie, Amy, and most importantly Sadie's PhD advisor Yiling Chen. Yiling, thank you for trusting our vision with Sadie and supporting this research; this was one of the most challenging research projects I have undertaken, and your support meant so much to me. I am so happy to know Sadie is getting advised by someone she can look up to and get inspired by.

In my years at Brown, I was lucky to spend time with a number of PhD students: Brandon Woodard, Enrique Areyan, Arjun Prakash, and Kevin Wang. Brandon, thank you for the memories and the experiences, I cannot wait to see the things you will do. Enrique, thank you for helping me transition to Brown. Arjun, thank you for the countless hours we have spent working together, I am excited to see the great places life takes you. Kevin, I only wish we could have spent more time, but I trust our paths will cross again.

There was also a very special PhD student: Lucy Qin. Lucy, my kind, beautiful, free, and brilliant partner in crime. I do not know how I would have done this PhD without you. You have been with me on my hardest and happiest days, from Fox Point, Providence, to Chinatown, New York, to Miami Beach, Miami. People like you, and a friendship like ours come once in a lifetime. You have been one of the greatest finds of this journey.

Beyond the support I have received from people I met during my PhD, I am eternally grateful for the support of three groups of people. Chief amongst them, my friends: Eric Deng, Roysi Eskenazi, Timur Emre Civan, Audrey Everist, Lisa Inzerillo, Sebastian Torero, Rebecca Vitenzon.

The second group of people are the high school teachers who allowed me to get to where I am: Hélène Jousselin, Ayşin Dülger, Connie Pierce, and Christopher Conley.

Last, and finally, I am so grateful for the administrators at Brown, who did their best to get us through the Kafkaesque academic bureaucracy and logistical challenges: Lauren Clarke, Jesse Polhemus, and John Tracey-Ursprung.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Historically, the most successful mathematical models—those that have found widespread applications across various fields—have been distinguished by two key attributes: 1) their broad applicability and 2) their ability to provide comprehensive mathematical characterizations. These models—including those examined in this thesis, such as variational inequalities, pseudo-games, and Markov pseudo-games—have helped us make sense of our reality, whether it be physics (e.g., fluid dynamics (Duvaut and Lions, 1976)), or the main object of study of this thesis: economics (e.g., asset and commodity pricing (Arrow and Debreu, 1954; Dafermos, 1990)).

While these models which found widespread applications across disciplines—most of which were developed in the early 20th century—have improved our understanding of the world, deriving actionable conclusions from them requires solving them. Initially, researchers sought to obtain closed-form solutions to their models, but as analytical solutions are intractable beyond very simple applications, with the introduction of computers in the 1960s, researchers turned their attention to the use of algorithms to instead obtain numerical solutions for their models (Scarf, 1967a).

Thus, at the dawn of the Information Age in the 1970s (Rifkin, 2011), applied mathematics and computer science researchers, equipped with their models and building on the foundations of computational complexity theory (Lax, 1989), set their sights on discovering algorithms of broad applicability—capable of solving a vast array of mathematical optimization frameworks. In particular, computer scientists began the search for a "holy grail" algorithm with two key characteristics (Fortnow, 2013): (1) broad applicability, meaning the ability to accurately solve a wide range of mathematical modeling frameworks, and (2) polynomial-time efficiency, ensuring that the algorithm halts with a solution in a number of computational steps polynomial in the input size. While some initial successes were achieved, it soon became apparent that such an algorithm might not exist at all (Garey and Johnson, 2002).

Indeed, in the decades following, beginning with the seminal work of Cook (1971) and later advanced by Papadimitriou (1994), a series of computational complexity results established that any polynomial-time algorithm capable of solving one mathematical modeling framework could, in principle, be adapted to solve another in polynomial time. While these findings established the theoretical existence of a universally applicable algorithm for all such frameworks, they also underscored a fundamental limitation: more than half a century after their introduction, no polynomial-time algorithm has been discovered for solving even a single one of these models—except in very special cases.

This conjecture, namely the lack of existence of a holy grail algorithm, presents a significant challenge to the use of mathematical models in solving real-world problems. Yet, as the demand for such models grows—whether for assessing the impact of climate change or optimizing public expenditure—researchers and practitioners are often forced to rely on algorithms that terminate within a reasonable timeframe, albeit at the cost of reduced accuracy. This trend is particularly concerning, as these models are applied in high-stakes domains where inaccuracies can lead to misleading or even detrimental outcomes (see, for instance, (Kim and Kim, 2003)). *This thesis takes this challenge as its starting point, aiming to equip practitioners—particularly in economics—with an algorithmic framework for understanding and addressing the computational complexities inherent in their modeling problems.*

1.2 Scope and Thesis

Three mathematical modeling frameworks will be the object of study of this thesis:

1. **Variational inequalities** (Lions and Stampacchia, 1967): A mathematical model of problems whose set of solutions can be posed as the solutions to an inequality involving a function.
2. **Pseudo-games** (Arrow and Debreu, 1954): A mathematical model of problems whose set of solutions can be posed as the solution of a (static) multiagent optimization problem.
3. **Stochastic pseudo-games**: A mathematical model of problems whose set of solutions can be posed as the solution of a multiagent *sequential optimization* problem.

The two key characteristic linking these three frameworks is that a solution to them is always guaranteed to exist (under mild assumptions), and the existence of solution is established via a non-constructive fixed point theorem such as Kakutani's (Glicksberg, 1952; Kakutani, 1941) or Brouwer's fixed point theorem. In more technical terms, the problem of computing a solution to these models belongs to a class of problems known as PPAD (Papadimitriou, 1994; Daskalakis et al., 2009; Chen and Deng, 2005). While a solution to these problems is guaranteed to exist, it has now become a widely upon agreed conjecture that there does not exist an algorithm that can solve problems in the PPAD class in polynomial time (Yannakakis, 2009).

In this thesis, these frameworks will be used to model the problem of computing a solution to three different models of economies, all of which belong a class of well-established economic models known as **general equilibrium models** or, equivalently, **infinite Walrasian economies** (i.e., highly abstract models of economies based on the demand and supply for a set of goods):[1]

[1]In this thesis, I use the "Walrasian economy" terminology to refer to economies with a finite set of commodities. However, Walrasian economies can in general have infinitely many commodities (Prescott and Lucas, 1972), as is the case with Radner economies, which consist of a possibly infinite set of commodities, as

1. **(Finite) Walrasian economies**: A highly abstract model of an economy based on the demand and supply for a *finite* set of commodities.
2. **Arrow-Debreu economies**: Finite Walrasian economies in which the demand and supply is explicitly generated by consumers and firms, respectively.
3. **Radner economies**: An infinite-dimensional generalization of Arrow-Debreu economies which explicitly incorporates time and uncertainty in consumers' and firms' decisions.

A hierarchy of the above general equilibrium models, as well as others which will be discussed in the sequel, is depicted in Figure 1.1.

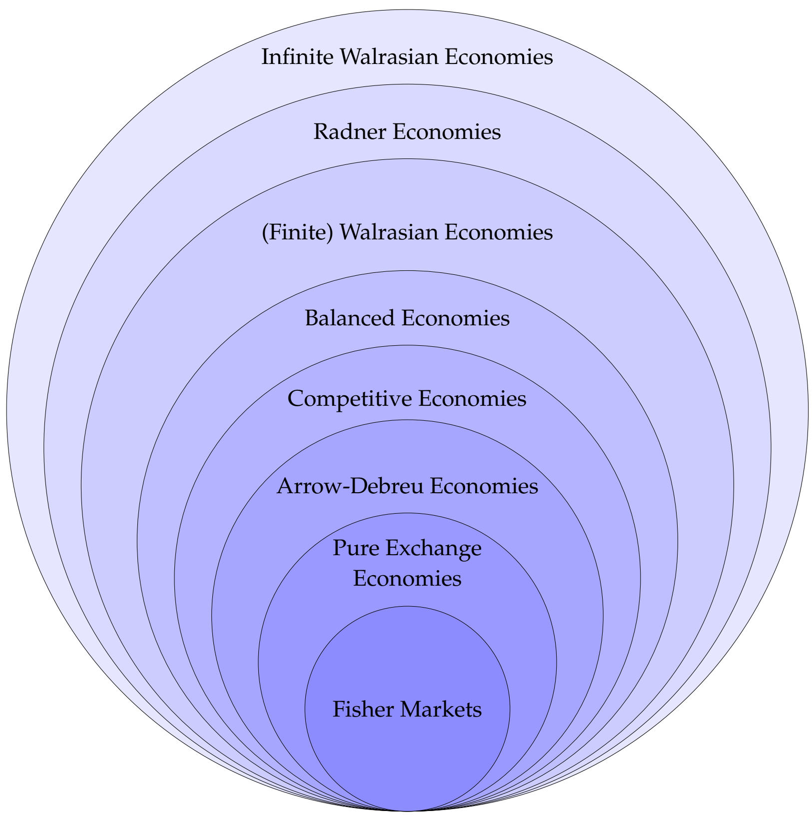

Figure 1.1: Hierarchy of the General Equilibrium Models Studied in this Thesis

General equilibrium models, and specifically the aforementioned models, are the foundation of much of mathematical economics, and are nowadays used in a myriad of impactful

the complete set of commodities of the economy is given by the union of all commodities across the potentially infinitely many states of the economy.

applications from resource allocation to public policy analysis. Unfortunately, the computation of a solution to these problems has also been shown to belong to the class of PPAD problems (Deng and Du, 2008; Chen and Teng, 2009), a class of problems that is unlikely to be solvable in polynomial time, except in special cases. As such, practitioners who use these models often resort to using algorithms that halt in a reasonable amount of time at the cost of increased inaccuracy. This practice raises concerns, as these algorithms are for instance used to solve models for public policy at central banks, with inaccurate solutions often leading to disastrous policy recommendations (see, for instance Kim and Kim (2003)).

In light of these discouraging facts, I will take an optimistic stance, and defend the following thesis.

Thesis

There exists a meaningful algorithmic theory of general equilibrium which allows practitioners to effectively trade-off accuracy and computational efficiency, and provides broadly-applicable algorithms which perform well in practice.

To argue this position, I will seek to answer the following twin questions:

1. Can we develop a unified mathematical and computational framework for solving a broad swath of general equilibrium models in a systematic way?

2. Can we develop broadly applicable algorithms which perform well in practice, and whose performance can be explained by this framework despite existing impossibility results?

My answer will be two-fold:

1. I will study the mathematical modeling frameworks of VIs, pseudo-games, and Markov pseudo-games, and develop algorithmic[2] methods with theoretical guarantees to solve these models.

2. I will provide characterizations of Walrasian, Arrow-Debreu, and Radner economies using the aforementioned mathematical modeling frameworks and use these characterizations devise broadly applicable computational methods that accurately and efficiently solve these general equilibrium models in practice, and provide a theoretical explanation of their strong empirical performance.

1.3 Outline

This thesis is organized in three major parts, which are presented following a review (Chapter 2) of the necessary mathematical background to understand its content. For readers familiar with the mathematical background, I include a "*too long; did not read*" section (Section 2.1), which summarizes the mathematical notation adopted in this thesis.

Each part of this thesis is broken down into two main chapters with the first chapter of each part describing an optimization framework, its mathematical and algorithmic properties, and the second chapter describing its application to a type of general equilibrium model.[3] I summarize the outline of these three major parts below.

[2]My use of the terms of "algorithmic" and "computational" is throughout the thesis mostly interchangeable; at points, however, I will use the "algorithmic" terminology to insist on the theoretical aspects of algorithms, while I will use "computational" to in addition refer to the empirical behavior of the algorithms.

[3]Part I, in addition to these two main chapters also includes an additional chapter on Fisher markets (Chapter 6) as a specific example of Walrasian economies, for illustrative purposes.

Outline of Major Thesis Parts

(Part I) Variational Inequalities and Walrasian Economies

(Chapter 4) Variational Inequalities

(Chapter 5) Walrasian Economies

(Part II) Pseudo-Games and Arrow-Debreu Economies

(Chapter 9) Pseudo-Games

(Chapter 10) Arrow-Debreu Economies

(Part III) Markov Pseudo-Games and Radner Economies

(Chapter 12) Markov Pseudo-Games

(Chapter 13) Radner Economies

Each chapter that concerns an optimization framework consists of three major sections: 1) mathematical background, 2) first-order methods to solve the optimization framework, 3) merit function methods to solve the optimization framework. The precise meaning "first-order method" and "merit function method" is described in the chapter relevant to each optimization framework.

Similarly, each chapter that concerns general equilibrium models consists of three major sections: 1) mathematical background and formulation within the optimization framework presented in the preceding chapter, 2) application of the first-order method presented in the preceding chapter, 3) application of the merit function method presented in the preceding chapter.

The organization of these parts in the given order serves three purposes. First, the results introduced in Chapter 4 are used in the rest of the thesis, thus requiring Part I to come before all others. Second, VIs (respectively, Walrasian economies) can model as a special case pseudo-games (respectively, Arrow-Debreu economies); thus, the results in Part I can provide additional insights about the results in Part II. Third, Markov pseudo-games (respectively, Radner economies) can be seen as infinite-dimensional (respectively, infinitely many commodity) generalizations of VIs and pseudo-games (respectively, Walrasian economies and Arrow-Debreu) economies, which the literature has started exploring only recently, and as such the results in Part I and Part II allow the reader to contextualize the results in Part III more effectively, and provides an open-ending for a new and exciting research direction on infinite-dimensional optimization and economies with infinitely many commodities.

1.4 Contributions

A high-level summary of the contributions of each chapter of the thesis can be found below. For a more detailed explanation of the contributions of each chapter, I refer the reader to the Contributions Section of the part in which the chapter is contained (Sections 3.3, 8.3, 11.3). Additionally, a summary of the main computational results for the optimization frameworks and general equilibrium models studied in this thesis can be found in Table 1.2 and Table 1.1, respectively.

(Chapter 4: Variational Inequalities): The mirror extragradient class of algorithms (Algorithm 3, Chapter 4) is introduced, with best-iterate polynomial-time convergence established for variational inequalities (VIs) that satisfy the Minty condition and are pathwise Bregman-continuous (e.g., Lipschitz-continuous VIs).[4] This result generalizes the extragradient method analysis of Huang and Zhang (2023) and extends the convergence guarantees of Zhang and Dai (2023) from unconstrained to constrained domains. Additionally, conditions for the *local* convergence of the mirror extragradient algorithm to an ε-strong solution in Bregman-continuous VIs *in the absence of* the Minty condition are established, representing the first known result of this kind. Finally, for general VIs, a polynomial-time

[4]Bregman-continuity is a generalization of Lipschitz-continuity in terms of the Bregman divergence, while pathwise Bregman-continuity is a further weakening of the Bregman-continuity condition, which requires Bregman-continuity to only hold over trajectories of the algorithm. See, Chapter 4 for additional definitions and explanation.

globally convergent class of merit function methods is developed to compute a solution satisfying the necessary conditions for a strong solution of the VI.

(CHAPTER 5: WALRASIAN ECONOMIES): A computationally tractable characterization of Walrasian equilibria in balanced economies is established as the strong solutions of a variational inequality (VI) satisfying the Minty condition within a unit box constraint. This result leads to the introduction of the mirror *extratâtonnement process* (Algorithm 6, Chapter 5), a novel price-adjustment method based on the mirror extragradient approach, whose polynomial-time convergence is proven in all balanced economies satisfying pathwise Bregman-continuity. Polynomial-time convergence is further demonstrated in competitive economies that are variationally stable under bounded excess demand elasticity, extending prior polynomial-time *tâtonnement* convergence results under the Gross Substitutes (GS), Weak Gross Substitutes (WGS), and Weak Axiom of Revealed Preferences (WARP) conditions (see, Figure 1.2 for additional details, and Chapter 5 for precise definitions) to a much larger class. The process is also shown to converge in polynomial time within the Scarf economy, marking the first such result for a natural discrete-time price-adjustment method. Experimental validation confirms the theoretical assumptions necessary for convergence and demonstrates efficient computation of Walrasian equilibria in large-scale competitive economies, including PPAD-complete cases such as Leontief economies, for which fast and reliable convergence is obtained. Finally, for general, potentially non-balanced Walrasian economies, a polynomial-time globally convergent class of merit function methods is developed to compute a solution satisfying the necessary conditions for a Walrasian equilibrium.

(CHAPTER 6: HOMOTHETIC FISHER MARKET): The maximum absolute value of the Hicksian price elasticity of demand is identified as a key parameter for analyzing the convergence of (entropic) *tâtonnement* in an important class of Walrasian economies known as homothetic Fisher markets. A sublinear convergence rate of $O(^{(1+\epsilon^2)}/_T)$ is established, where ϵ represents the maximum Hicksian price elasticity of demand across buyers. This result generalizes existing convergence analyses for CES and nested CES utilities, unifying previously disjointed convergence and non-convergence findings. It encompasses the full spectrum of (nested) CES utilities, including Leontief and linear utilities, recovering the best-known rate of $O(1/T)$ for Leontief markets ($\epsilon = 0$) and confirming the non-convergent behavior of *tâtonnement* in linear markets as $\epsilon \to \infty$. Known existing computational results for the convergence of *tâtonnement* in Fisher markets in light of this result is summarized in Figure 3.1a.

(CHAPTER 9: PSEUDO-GAMES) The existence of variational equilibrium in quasiconcave pseudo-games with jointly convex constraints is first re-established, followed by the introduction of first-order variational equilibrium, which is shown to exist in a broader class of pseudo-games than previously known—specifically, smooth games with jointly convex constraints. An equivalence is then established between (first-order) variational equilibria in pseudo-games and strong solutions of variational inequalities, leading to the characterization of a new class of pseudo-games, termed variationally stable pseudo-games with jointly convex constraints. For this class, first-order variational equilibrium can be computed in polynomial time using a novel learning dynamic called the mirror extragradient learning dynamics (Algorithm 7, Chapter 9). In the special case where the pseudo-game is also concave, this result extends to variational equilibrium computation, representing the broadest known result of its kind. Finally, for general pseudo-games with jointly convex constraints that are not necessarily variationally stable, a polynomial-time globally convergent class of merit function methods is developed to compute a solution that satisfies the necessary conditions for variational equilibrium.

(CHAPTER 10: ARROW-DEBREU ECONOMIES): Novel mathematical characterizations of Arrow-Debreu equilibrium in Arrow-Debreu economies are developed. First, it is re-established that the set of Arrow-Debreu equilibria in any quasiconcave Arrow-Debreu economy corresponds to the set of generalized Nash equilibria of the Arrow-Debreu pseudo-game. Due to the intractability of this characterization, a new formulation is introduced, expressing the set of Arrow-Debreu equilibria in any concave pure exchange economy as the set of generalized Nash equilibria of the trading post pseudo-game, a variationally stable pseudo-game with jointly convex constraints. The mirror extragradient learning dynamics are then applied to this pseudo-game, yielding a market dynamic termed the mirror extratrade dynamic. While the trading post pseudo-game is not concave, it is shown to be pseudoconcave, allowing an approximate first-order variational equilibrium to be computed in polynomial time, with asymptotic convergence to a variational equilibrium—and thus to an Arrow-Debreu equilibrium of the associated concave pure exchange economy—to the best of my knowledge the broadest convergence result of its kind. Finally, for general, potentially non-concave Arrow-Debreu economies, a polynomial-time globally convergent class of merit function methods is developed to compute a solution satisfying the necessary conditions for an Arrow-Debreu equilibrium.

(CHAPTER 12: MARKOV PSEUDO-GAMES): Markov pseudo-games are introduced as a generalization of Markov games (i.e., games with time and uncertainty), where other players' actions influence both rewards and available actions. The existence of pure generalized Markov perfect equilibria (GMPE) is established in concave Markov pseudo-games, extending to the dynamic stochastic setting, Arrow-Debreu's equilibrium existence results for (static) concave pseudo-games. This also implies the existence of pure Markov perfect equilibria in a broad class of continuous-action Markov games, where previously only mixed-strategy equilibria were known to exist (Fink, 1964; Takahashi, 1964). Although computing GMPE is PPAD-hard in general, the problem is reformulated as a generative adversarial learning task, where a generator proposes an equilibrium policy profile, and an adversary produces best responses. Leveraging advances in generative adversarial learning, it is shown that under mild assumptions, a policy profile that is a stationary point of exploitability (players' cumulative maximum regret) can be computed in polynomial time. This result applies to Markov pseudo-games with a bounded best-response mismatch coefficient, which requires that states explored by any GMPE are sampled sufficiently often under the initial state distribution. This approach parallels known computational results for zero-sum Markov games. As these theoretical guarantees hold for policies represented by neural networks, this provides the first deep reinforcement learning algorithm with theoretical guarantees for general-sum games—a very broad class of general-sum games, at that.

(Chapter 13: Radner Economies): An extension of Magill and Quinzii's infinite horizon exchange economy (Magill and Quinzii, 1994), termed the Radner economy, is introduced, generalizing the model to arbitrary assets while restricting the transition dynamics to be Markovian. This restriction enables a proof of existence for a recursive Radner equilibrium (RRE), a Radner equilibrium independent of the initial state distribution, simplifying the equilibrium policy domain to the space of states rather than histories and making computation more tractable. The set of RREs in any Radner economy is reformulated as the set of generalized Markov perfect equilibria (GMPE) of an associated Markov pseudo-game, extending prior results that were limited to economies with a single consumer, commodity, or asset. This formulation further implies that a stationary point of exploitability in the associated Markov pseudo-game can be computed in polynomial time. To validate these theoretical results, a the method is implemented as a generative adversarial policy network and applied to three Radner economies with distinct utility functions. Experimental findings indicate that the method produces approximate equilibrium policies that are much closer to GMPE than those generated by a standard macroeconomic baseline for solving stochastic economies.

1.5 Historical and Academic Context

1.5.1 General Equilibrium Theory: The Foundations of Economic Modeling

Key to the historical development of mathematical models in economics was a need to understand how economies functioned (i.e., the emergence of demand, supply, and prices), dating back to as early as the 18th century to French-Irish economist Richard Cantillon's

Part	Chapter	Economy Class	Economy Subclass	Solution Concept	Algorithm	Convergence Rate	Result Reference
I	5	Walrasian Economies	Balanced + Pathwise Bregman-Continuous	Walrasian Equilibrium	Mirror Extratâtonnement	$O(1/\sqrt{\tau})$	Theorem 5.4.1
I	5	Walrasian Economies	Competitive + Variationally Stable + Elastic + Bounded	Walrasian Equilibrium	Mirror Extratâtonnement	$O(1/\sqrt{\tau})$	Theorem 5.4.2
I	5	Walrasian Economies	Scarf Economy	Walrasian Equilibrium	Mirror Extratâtonnement	$O(1/\sqrt{\tau})$	Corollary 5.4.4
I	5	Walrasian Economies	Lipschitz-Continuous + Lipschitz-Smooth	Stationary Point of Walrasian Merit Function	Mirror Potential	$O(1/\tau)$	Theorem 5.5.1
I	6	Walrasian Economies	Homothetic Fisher Markets + Bounded Hicksian Demand Elasticity	Walrasian Equilibrium	Entropic Tâtonnement	$O(1/\tau^2)$	Theorem 6.5.1
II	10	Arrow-Debreu Economies	Concave Pure Exchange Economy + Lipschitz-Smooth Utilities	Arrow-Debreu Equilibrium	Mirror Extratrade Dynamics	$O(1/\sqrt{\tau})$	Theorem 10.3.1
II	10	Arrow-Debreu Economies	Quasiconcave + Lipschitz-Smooth Utilities	Stationary Point of Arrow-Debreu Exploitability	REDA	$O(1/\tau)$	Theorem 10.4.1
II	10	Arrow-Debreu Economies	Lipschitz-Smooth Utilities + Lipschitz-Smooth Utility Gradients	Stationary Point of Arrow-Debreu Variational Exploitability	Mirror Variational Dynamics	$O(1/\tau)$	Theorem 10.4.2
III	13	Radner Economies	Concave + Lipschitz-Smooth Utilities	Stationary Point of Radner (State) Exploitability	TTSGDA	$O(1/\sqrt[6]{\tau})$	Theorem 10.4.1

Table 1.1: Summary of Main Computational Results for General Equilibrium Models

Part	Chapter	Optimization Framework	Class of Framework	Solution Concept	Algorithm	Convergence Rate	Result Reference
I	4	Variational Inequalities	Minty + Pathwise Bregman-Continuous	Strong (Stampacchia) Solution	Mirror Extragradient Algorithm	$O(1/\sqrt{\tau})$	Theorem 4.3.1
I	4	Variational Inequalities	"Local" Minty + Bregman-Continuous	Strong (Stampacchia) Solution	Mirror Extragradient Algorithm	$O(1/\sqrt{\tau})$ Local Convergence	Theorem 4.3.2
I	4	Variational Inequalities	Lipschitz-Continuous + Lipschitz-Smooth	Stationary Point of Regularized Primal Gap	Mirror Potential Algorithm	$O(1/\tau)$	Theorem 4.4.1
II	9	Pseudo-Games	Concave + Variationally Stable + Lipschitz-Smooth + Jointly Convex Constraints	Variational Equilibrium	Mirror Extragradient Dynamics	$O(1/\sqrt{\tau})$	Theorem 9.4.1
II	9	Pseudo-Games	Variationally Stable + Lipschitz-Smooth + Jointly Convex Constraints	First-Order Variational Equilibrium	Mirror Extragradient Dynamics	$O(1/\sqrt{\tau})$	Theorem 9.6.1
II	9	Pseudo-Games	Concave + Lipschitz-Smooth + Jointly Convex Constraints	Stationary Point of Regularized Exploitability	REDA	$O(1/\tau)$	Theorem 9.4.3
II	9	Pseudo-Games	Lipschitz-Smooth Payoff + Lipschitz-Smooth Payoff Gradients + Jointly Convex Constraints	Stationary Point of Variational Exploitability	Mirror Variational Dynamics	$O(1/\tau)$	Theorem 9.6.2
III	12	Markov Pseudo-Games	Concave + Lipschitz-Smooth	Stationary Point of (State) Exploitability	TTSGDA	$O(1/\sqrt[6]{\tau})$	Theorem 12.3.1

Table 1.2: Summary of Main Computational Results for Optimization Frameworks

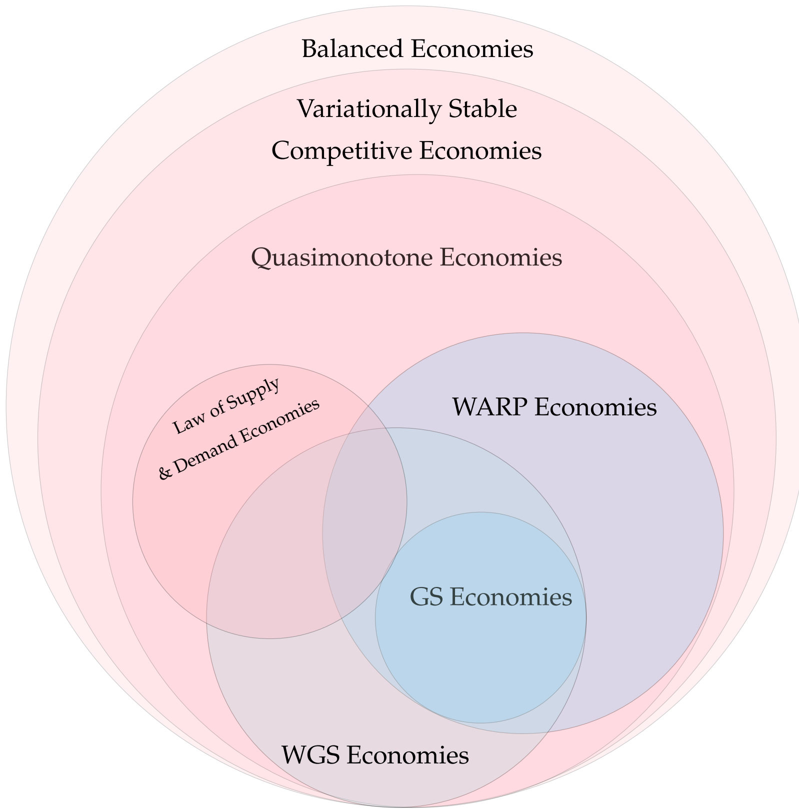

Figure 1.2: Walrasian Economies for which there exists a polynomial-time price-adjustment process. Economies for which I prove polynomial-time convergence of mirror *extratâtonnement* are depicted in pink hues (i.e., balanced, variationally stable competitive, quasimonotone, law of supply and demand economies), while economies for which polynomial-time convergence was known are depicted in blue hues (i.e., WGS, GS, WARP economies). The convergence result for balanced economies holds under the assumption of pathwise Bregman-continuity, whose plausibility I verify in experiments (Section 5.4.3, Chapter 5, Part I), while the convergence result for variational stable competitive economies holds under the assumption that the price elasticity of excess demand is bounded (see, Chapter 5, Part I for additional details). The polynomial-time price-adjustment process for economies in blue is *tâtonnement*. The convergence result for WARP economies was introduced by Uzawa (1960); the first weakly polynomial-time convergence result (i.e., without any elasticity boundedness assumptions) for WGS economies was introduced by Codenotti et al. (2005), with Cole and Fleischer (2008) proving a strongly polynomial-time convergence result (i.e., in terms of elasticity bounds).

work (Cantillon, 1755). While a great number of economists including Adam Smith (Smith, 1937), David Ricardo (Ricardo, 1895), John Stuart Mill (Mill, 1965), and Alfred Marshall (Marshall, 1910) would make great contributions to our understanding of demand, supply, and prices, it would not be until the pioneering work of French economist Léon Walras (Walras, 1896) that a clear modeling paradigm for economies would emerge.

Walras formulated a mathematical model of markets (nowadays known as a **Walrasian market**) as a system of resource allocation comprising supply and demand functions that

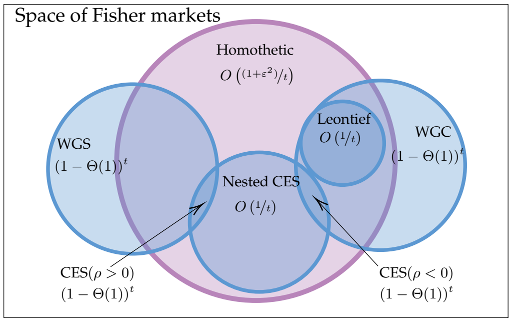

(a) The convergence rates of *tâtonnement* for different Fisher markets. We color previous contributions in blue, and our contribution in red, i.e., we study homothetic Fisher markets where ϵ is the maximum absolute value of the price elasticity of Hicksian demand across all buyers. We note that the convergence rate for WGS markets does not apply to markets where the price elasticity of Marshallian demand is unbounded, e.g., linear Fisher markets; likewise, the convergence rate for nested CES Fisher markets does not apply to linear or Leontief Fisher markets.

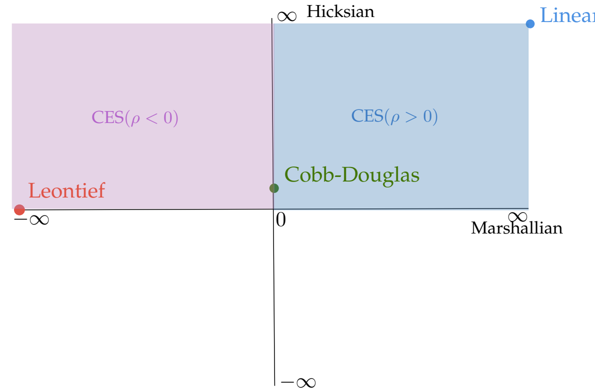

(b) Cross-price elasticity taxonomy of well-known homogeneous utility functions. There are no previously studied utility functions in the space of utility functions with negative Hicksian cross-price elasticity. Future work could investigate this space and prove faster convergence rates than those provided in this thesis. We note that our convergence result covers the entire spectrum of this taxonomy (excluding the limits of the y-axis).

Figure 1.3: A summary of known results in Fisher markets.

map values for resources, called **prices**, to quantities of resources—*ceteris paribus*, i.e., all else being equal. Walras also defined a steady state of a market, which he called **competitive** (nowadays, also called a **Walrasian**) **equilibrium**, as prices s.t. the demand is **feasible**, i.e., the demand for each resource is less than or equal to its supply, and **Walras' law** holds, i.e., the value of the supply is equal to the value of the demand. Unlike in Walras' model, in real-world markets, all else is not equal, and markets do not exist in isolation but are part of an **economy**. Indeed, the supply and demand of resources in one market depend not only on prices in that market, but also on the supply and demand of resources in other markets. If every market in an economy is simultaneously at a competitive equilibrium, Walras' law holds for the economy as a whole; this steady state, now a property of the economy, is called a **general equilibrium**.

As such, Walras' early forays in economic modeling would leave two important questions open. First, it would be unclear if Walras' market model could model an entire economy with consumers and firms and a number of different markets. Second, Walras did not provide conditions that guarantee the existence of a competitive equilibrium and it was not clear if such prices existed. The question of whether Walras' model could be seen as a model of an economy, and under what conditions competitive equilibrium prices exist, would remain open for over half a century until Arrow and Debreu's seminal **model of a competitive economy** (nowadays also called **Arrow-Debreu economies**) for which Arrow and Debreu (1954) proved the existence of a competitive equilibrium.

Remark 1.5.1 [Some history]**.**
Following the second world war, European and American governments had to preside over the reconstruction of devastated economies. This required an improved understanding of the role of public expenditure in economic activity, since the war had led to an unprecedented increase in the role of government in the economy, a change that unsustainable. While in 1944, the American government's spending at all levels accounted for 55 percent of gross domestic product (GDP), by 1947, government spending had dropped 75 percent in real terms, or from 55 percent of GDP to just over 16 percent of GDP (Bureau of Economic Analysis, 2021). Fortunately, when the Second World War (WW2) broke out, many European academics escaping the war moved to the United States, aiding the development of rigorous mathematical models of economies, which would be key to understanding how to manage the role of public expenditure in economic growth. These academics concentrated primarily at the University of Chicago, where the Cowles Commission was founded in 1939, with the Viennese maxim, "Science is Measurement" (Mitra-Kahn, 2005). The Cowles commission aimed to link mathematics and economics (Mitra-Kahn, 2005), and played a crucial role in the development of mathematical microeconomic models that have become the foundation of modern economics. These efforts, initiated by the Cowles Commission, culminated in the seminal work of Keneth Arrow and Gérard Debreu, who proved the existence of general equilibria in a very general setting (Arrow and Debreu, 1954).

In their model, Arrow and Debreu posit a set of resources, modeled as commodities, each of which is assigned a price; a set of consumers, each choosing a quantity of each commodity to consume that maximizes their utility function in exchange for their endowment (e.g., labor); and a set of firms, each choosing a profit-maximizing quantity of each commodity to produce, with prices determining **aggregate demand**, i.e., the sum of the utility-maximizing consumptions across all consumers, and **aggregate supply**, i.e., the sum of endowments and profit-maximizing productions across all consumers and firms, respectively. Notably, Arrow and Debreu's model could be seen as a special case of Walras model defined by the aggregate demand and aggregate supply, and its solution being defined as the associated competitive, now better called general, equilibrium. This reduction would demonstrate that Walras' much simpler model can accurately model an economy without the need to assume that all else is equal, allowing us to call his model a **Walrasian economy** rather than a market *de jure*. Even more importantly, under mild assumptions Arrow-Debreu proved that a general equilibrium exists in their model, hence providing sufficient and meaningful conditions for the existence of a competitive equilibrium in Walrasian economies.

Nevertheless, soon after Arrow and Debreu's monumental achievement, another issue would emerge. Arrow-Debreu and Walrasian economies are static economies, which do not explicitly model time and uncertainty. That is, unlike real-world economies, all trade in such economies happens in a single time period when the world is at a given state. To get around this issue, Arrow and Debreu argued that commodities were to be seen as state and time contingent, with each one representing a good or service which can be bought or sold in a single time period, but that encodes delivery opportunities at a finite number of distinct points in a larger space that incorporates state and time. However, as this explanation would not be a realistic explanation of real-world economies, in the decades to come, economists would seek an answer to the question of whether Arrow-Debreu and Walrasian economies could represent an economy with time and uncertainty, and if a general equilibrium would be guaranteed to exist in economies with time and uncertainty.

The question would mostly remain open for nearly 20 years until Radner's introduction of the **stochastic competitive economy** (nowadays also called the **Radner economy** in which he proved the existence of a general equilibrium under suitable assumptions (Radner, 1972). The Radner economy, initialized at a state of the world, is a finite-horizon economy comprising a sequence of **spot markets** in which consumers and firms can purchase and sell commodities for immediate delivery, all linked across time by **asset markets** in which consumers and firms can buy or sell assets that deliver a payment when a particular state of the world occurs, with the economy stochastically transitioning to one of many other world states once consumers and firms have made their purchases. Commodities (respectively, assets) are assigned state and time contingent prices, which determine their **aggregate state and time contingent demand**, i.e., the sum of utility maximizing consumptions (respectively, asset portfolios), and **aggregate state and time contingent supply**, i.e., the sum of endowments and profit-maximizing productions (respectively, asset portfolios) across all consumers and firms, respectively. Similar to Arrow-Debreu economies, any Radner economy can be cast as a Walrasian economy given by the aggregate state and time contingent demand and supply, with its solution being defined as a general equilibrium of this Walrasian economy. Further, under suitable assumptions on the asset market, Arrow (1964) shows that any Radner economy can also be represented as an Arrow-Debreu economy, thus further demonstrating that Walrasian economies and their solution concept, the general equilibrium, can effectively model economies with time and uncertainty.

Nearly half a century later, Arrow-Debreu and Radner economies have become foundational pillars of modern mathematical economics, providing an explanation of the most important facets of any economy. While these models only scratch the surface of the mathematical models of economies developed since Arrow and Debreu's work, all such models share one common characteristic: they can be cast as Walrasian economies, with their solution corresponding to a general equilibrium of the associated Walrasian economy, thus leading to them colloquially called **general equilibrium models**.

In developing their models and establishing their existence results, Arrow and Debreu and Radner would pioneer the development of the theory of games and stochastic games, which would continue to play a key role in the development of various other branches of mathematical economics, such as mechanism design and financial economics. While the algorithmic theory developed in this thesis will be applied specifically to the analysis of algorithms for general equilibrium models, all the results developed within this thesis will be developed within broader and more abstract mathematical optimization frameworks, some of which were first studied by Arrow and Debreu (e.g., pseudo-games) and Radner (e.g., stochastic pseudo-games), and others developed subsequently (e.g., variational inequalities). As such, the algorithms and analyses I provide in this thesis are relevant not only to general equilibrium models, but also to other areas of mathematical economics such as the aforementioned areas of mechanism design and financial economics.

1.5.2 General Equilibrium Theory at the Origin of Mechanism Design

Much of mathematical microeconomic theory since the 1970s would focus on the development of mechanism design (i.e., mathematical and algorithmic frameworks for the design of markets), which has nowadays become a cornerstone of mathematical economics. In particular, the computational literature on economics has dedicated a great deal of resources, often at the expense of general equilibrium theory, to the development of an algorithmic theory of mechanism design, perhaps due to the financial incentives provided by the emergence of online market places. The goal of this digression is to clarify the historical connections between mechanism design and general equilibrium theory, and underscore the importance of developing the algorithmic theory of general equilibrium to further our understanding of algorithmic mechanism design.

Recall that the key political driver behind the development of general equilibrium theory was the need to better understand economies in order to optimally reduce public expenditure, which had drastically increased during World War II (see Remark 1.5.1). As part of this development, in separate papers, Arrow and Debreu independently but simultaneously proved the first and second welfare theorems of economics, which stated that 1) consumptions and productions associated with a general equilibrium are Pareto-efficient and 2) any collection of Pareto-optimal consumptions and productions can be associated with a general equilibrium (price system) (Arrow, 1951a; Debreu, 1951b).

Arrow and Debreu's results implied that competitive economies[5] are thus an efficient way to allocate resources, since they result in a Pareto-optimal distribution of resources—an inference which is contingent on the global stability of general equilibria, i.e., free markets actually settling into a general equilibrium. Unfortunately, Arrow and Debreu's results provided ambiguous conclusions on how to transition away from a war economy with high public expenditure. On one hand, Arrow and Debreu's results suggested that a post-war economy with no governmental intervention and no public expenditure was optimal, as free markets are a Pareto-efficient mechanism of resource allocation. On the other hand, real-world markets would never truly be free and it seemed like decreasing public expenditure significantly could be a disastrous economic choice as Paul Samuelson, 1970 Nobel Prize laureate, wrote in 1943: "some ten million men will be thrown on the labor market" (Mitra-Kahn, 2005), warning that it would be "the greatest period of unemployment and industrial dislocation which any economy has ever faced" (Mitra-Kahn, 2005).

The issue of public expenditure in a post-war world would become a central theme in Samuelson's research, who would provide the first rigorous definition of public goods. Based on his definition Samuelson would then derive what came to be known as the Samuelson condition: the first-order optimality condition associated with the optimal provision of public goods in terms of the demand and supply for public goods (Samuelson, 1954). Samuelson's work was groundbreaking in the sense that it moved the problem of allocating public spending from the realm of political theory to the realm of general equilibrium theory. Samuelson's analysis would be studied more rigorously in subsequent general equilibrium models, eventually culminating in a generalization of Arrow-Debreu's

[5] A competitive economy in this sense is one that resembles an Arrow-Debreu economy.

model of a competitive economy called the private ownership economy with public goods, a model which differentiates between public and private goods (Mitra-Kahn, 2005).

An important result that emerged from this line of work started by Samuelson, which was proven by Foley, is that a general equilibrium (also known as a **Lindahl-Foley equilibrium**) exists in a general equilibrium model with private ownership and public goods (also known as a **Lindahl-Foley economy**), and that the first and second welfare theorems apply (Foley, 1970). Foley's results confirmed the role that governments have been attributed to achieving, namely a Pareto-optimal allocation of resources via redistribution policies (Mitra-Kahn, 2005). Although a gross simplification of the conclusion, this meant that governments could direct their public expenditures using the first-order optimality conditions determined by Samuelson, which now could be interpreted as the welfare-maximizing quantities of public goods prescribed by the general equilibrium of the Lindahl-Foley economy, to provide optimal levels of public goods, ensuring their economies settle into a Pareto-optimal allocation of resources!

While a Lindahl-Foley equilibrium provides us with a way to determine an optimal quantity of public goods to provide, computing such an optimal provision of public goods is not straightforward, because it requires the policy maker to know consumers' true preferences. One possible workaround is to elicit consumers' preferences; however, as it turns out, consumers have an incentive to lie about their preferences over public goods, as they can obtain a more favorable allocation of public goods prescribed by the Lindahl-Foley equilibrium by doing so. This in turn might lead the policy maker to compute a non-optimal provision of public goods.

The issue of incentive-compatibility, i.e., consumers not reporting their preferences truthfully, in the provision of public goods led to the development of mechanism design. In order to ensure that economies achieve a Pareto-optimal allocation of resources, governments had to choose their equilibrium levels of public expenditure, yet to compute these equilibrium levels, governments had to be able to elicit the true preference of consumers over public goods. This required the development of the mechanism design literature, which would provide a new formalization of social and economic interactions that accounted for incentives. The consistent objective of the mechanism design literature, which was introduced in seminal papers by Hurwicz in the 1960s and 70s (Hurwicz, 1972; 1960; 1979), was in fact to provide a more powerful framework than general equilibrium theory to address the issue of incentive compatibility. From this vantage point, mechanism design can be seen as a generalization of general equilibrium theory. More precisely, the general equilibrium or the competitive equilibrium is a particular outcome of a game, just like the VCG outcome is, and the Walrasian mechanism, i.e., the function which takes as input the preferences of consumers, computes and outputs a competitive equilibrium, is an instance of a mechanism just like the VCG mechanism is.

More importantly, however, history tells us that without general equilibrium theory, there would be no mechanism design. Perhaps an analogy is fitting here: general equilibrium is to mechanism design, as the normal distribution is to probability theory. Just like one cannot imagine a theory of probability without a solid understanding of the statistical and algorithmic properties of the normal distribution, one cannot expect a proper understanding of mechanism design without a proper understanding of the economic and algorithmic properties of general equilibria. History is the biggest testament to this statement: without general equilibrium theory's incentive issue, Hurwicz and others would not have been inspired to develop the theory of mechanism design. As such, the algorithms and analyses introduced in this thesis inform not only general equilibrium theory but also other branches of mathematical economics, perhaps most notably, mechanism design.

Chapter 2

Mathematical Background

2.1 Too Long; Did Not Read

For the reader familiar with the mathematical background, in this section, I provide a brief overview of the notational conventions and mathematical definitions used throughout this thesis. For the remainder of this thesis, excluding the conclusion, I will use the pronoun "we" instead of "I" to maintain narrative fluidity in the mathematical exposition.

2.1.1 Notation

We adopt the following calligraphic conventions to insist on the nature of the mathematical object at hand: We use calligraphic uppercase letters to denote sets (e.g., $\mathcal{X}$), bold uppercase letters to denote matrices (e.g., $\boldsymbol{X}$), bold lowercase letters to denote vectors (e.g., $\boldsymbol{p}$), lowercase letters to denote scalar quantities (e.g., x), and uppercase letters to denote random variables (e.g., X). We denote the ith row vector of a matrix (e.g., $\boldsymbol{X}$) by the corresponding bold lowercase letter with subscript i (e.g., $\boldsymbol{x}_i$). Similarly, we denote the jth entry of a vector (e.g., $\boldsymbol{p}$ or $\boldsymbol{x}_i$) by the corresponding lowercase letter with subscript j (e.g., p_j or x_{ij}). We denote functions by a letter determined by the value of the function, e.g., f if the mapping is scalar valued, $\boldsymbol{f}$ if the mapping is vector valued, and $\mathcal{F}$ if the mapping is set valued.

For any collection of sets $\{\mathcal{A}_i\}_{i \in [n]}$, we define the notation $(\boldsymbol{a}_i, \boldsymbol{a}_{-i}) \doteq (\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n) \in \bigtimes_{i \in [n]} \mathcal{A}_i$, where $\boldsymbol{a}_{-i} \in \bigtimes_{i' \in [n], i' \neq i} \mathcal{A}_{i'}$ denotes $(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_{i-1}, \boldsymbol{a}_{i+1}, \ldots, \boldsymbol{a}_n)$ with the ith entry $\boldsymbol{a}_i \in \mathcal{A}_i$ removed.

We denote the set $\{1, \ldots, n\}$ by $[n]$, the set $\{n, n+1, \ldots, m\}$ by $[n : m]$, the set of natural numbers by $\mathbb{N}$, and the set of real numbers by $\mathbb{R}$. We denote the positive and strictly positive elements of a set using a $+$ or $++$ subscript, respectively, e.g., $\mathbb{R}_+$ and $\mathbb{R}_{++}$. For any $n \in \mathbb{N}$, we denote the n-dimensional vector of zeros and ones by $\mathbf{0}_n$ and $\mathbf{1}_n$, respectively.

2.1.2 Mathematical Definitions

We let $\Delta_n = \{\boldsymbol{x} \in \mathbb{R}^n_+ \mid \sum_{i=1}^n x_i = 1\}$ denote the unit simplex in $\mathbb{R}^n$, and $\Delta(A)$ denote the set of all probability measures over a given set A. We also define the support of a probability density function $f \in \Delta(\mathcal{X})$ as $\operatorname{supp}(f) \doteq \{\boldsymbol{x} \in \mathcal{X} \mid f(\boldsymbol{x}) > 0\}$. Finally, we denote the orthogonal projection operator onto a set C by Π_C, i.e., $\Pi_C(\boldsymbol{x}) \doteq \arg\min_{\boldsymbol{y} \in C} \|\boldsymbol{x} - \boldsymbol{y}\|^2$.

For any $\varepsilon \geq 0$, we write $\mathcal{B}_\varepsilon[x] = \{x' \in \mathcal{M} \mid d(x, x') \leq \varepsilon\}$ and $\mathcal{B}_\varepsilon(x) = \{x' \in \mathcal{M} \mid d(x, x') < \varepsilon\}$ to denote the closed and open ε-ball centered at $x \in \mathcal{M}$, respectively.

For any real number $a \in \mathbb{R}$, $a\mathcal{X}$ denotes the (Minkowski) product, i.e., $a\mathcal{Y} \doteq \{ax \mid x \in \mathcal{X}\}$; $\mathcal{X} + \mathcal{Y}$ denotes the (Minkowski) sum of $\mathcal{X}$ and $\mathcal{Y}$, i.e., $\mathcal{X} + \mathcal{Y} \doteq \{x + y \mid x \in \mathcal{X}, y \in \mathcal{Y}\}$; and $\mathcal{X} - \mathcal{Y}$ denotes the (Minkowski) difference of $\mathcal{X}$ and $\mathcal{Y}$, i.e., $\mathcal{X} - \mathcal{Y} \doteq \{x - y \mid x \in \mathcal{X}, y \in \mathcal{Y}\}$.

We denote by $\mathbb{1}_{\mathcal{C}}(\boldsymbol{x})$ the indicator function of a set $\mathcal{C}$, with value 1 if $\boldsymbol{x} \in \mathcal{C}$ and 0 otherwise.

Given two vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$, we write $\boldsymbol{x} \geq \boldsymbol{y}$ or $\boldsymbol{x} > \boldsymbol{y}$ to mean component-wise $\geq$ or $>$, respectively.

For any set $\mathcal{C}$, we denote the diameter by $\operatorname{diam}(\mathcal{C}) \doteq \max_{\boldsymbol{c}, \boldsymbol{c}' \in \mathcal{C}} \|\boldsymbol{c} - \boldsymbol{c}'\|$.

We define the gradient operator $\nabla_{\boldsymbol{x}}$ as the operator which takes as input a function $f : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$, and outputs a vector-valued function consisting of the partial derivatives of f w.r.t. $\boldsymbol{x}$. We denote the derivative operator (respectively, partial derivative operator w.r.t. $\boldsymbol{x}$) of any function $\boldsymbol{g} : \mathcal{X} \times \mathcal{Y} \to \mathcal{Z}$ by $\partial \boldsymbol{g}$ (respectively, $\partial_{\boldsymbol{x}} \boldsymbol{g}$). We define the subdifferential of any function $f : \mathcal{X} \times \mathcal{Y} \to \mathbb{R}$ w.r.t. variable $\boldsymbol{x}$ at a point $(\boldsymbol{a}, \boldsymbol{b}) \in \mathcal{X} \times \mathcal{Y}$ by $\mathcal{D}_{\boldsymbol{x}} f(\boldsymbol{a}, \boldsymbol{b}) \doteq \{\boldsymbol{h} \mid f(\boldsymbol{x}, \boldsymbol{b}) \geq f(\boldsymbol{a}, \boldsymbol{b}) + \boldsymbol{h}^T(\boldsymbol{x} - \boldsymbol{a})\}$,

Functions. Given a Euclidean vector space $\mathcal{X} \subseteq \mathbb{R}^n$, we define its dual space $\mathcal{X}^*$ as the set of all linear maps $\boldsymbol{f} : \mathcal{X} \to \mathbb{R}^n$. Let $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$, $(\mathcal{Y}, \|\cdot\|_{\mathcal{Y}})$ be normed vector spaces. Consider a function $\boldsymbol{f} : \mathcal{X} \to \mathcal{Y}$. $\boldsymbol{f}$ is **continuous** if for all sequences $\{x^{(n)}\}_{n \in \mathbb{N}}$ s.t. $x^{(n)} \to x \in \mathcal{X}$, it holds that $f(x^{(n)}) \to f(x)$. Given $\ell \geq 0$, $\boldsymbol{f}$ is ℓ-**Lipschitz continuous** on $\mathcal{A} \subseteq \mathcal{X}$ iff for all $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathcal{A}$, $\|f(\boldsymbol{x}_1) - f(\boldsymbol{x}_2)\|_{\mathcal{Y}} \leq \ell \, \|\boldsymbol{x}_1 - \boldsymbol{x}_2\|_{\mathcal{X}}$. Consider a function $f : \mathcal{X} \to \mathbb{R}$. f is **convex** iff for all $\lambda \in [0, 1]$ and $\boldsymbol{x}, \boldsymbol{x}' \in \mathcal{X}$, $f(\lambda \boldsymbol{x} + (1 - \lambda)\boldsymbol{x}') \leq \lambda f(\boldsymbol{x}) + (1 - \lambda) f(\boldsymbol{x}')$. Given $\mu \geq 0$, f is μ-strongly-convex, iff $\boldsymbol{x} \mapsto f(\boldsymbol{x}) - {}^{\mu}\!/\!{}_{2}\|\boldsymbol{x}\|^2$ is convex.

Correspondences. Let $(\mathcal{X}, \langle \cdot, \cdot \rangle)$ be an inner product space. Consider a correspondence $\mathcal{R} : \mathcal{X} \rightrightarrows \mathcal{X}^*$. $\mathcal{R}$ is **continuous** if for any sequence $\{x^{(n)}\}_{n \in \mathbb{N}_+} \subset \mathcal{X}$ such that $x^{(n)} \to x$, it holds that $\mathcal{R}(x^{(n)}) \to \mathcal{R}(x)$. $\mathcal{R}$ is **upper hemicontinuous** if for any sequence $\{(x^{(n)}, y^{(n)}\}_{n \in \mathbb{N}_+} \subset \mathcal{X} \times \mathcal{X}^*$ such that $(x^{(n)}, y^{(n)}) \to (x, y)$ and $y^{(n)} \in \mathcal{R}(x^{(n)})$ for all $n \in \mathbb{N}_+$, it also holds that $y \in \mathcal{R}(x)$. $\mathcal{R}$ is **closed-valued** (resp. **compact-valued** / **convex-valued** / **singleton-valued**) iff for all $x \in \mathcal{X}$, $\mathcal{R}(x)$ is closed (resp. compact / convex / a singleton). $\mathcal{R}$ is **monotone** iff for all $x, x' \in \mathcal{X}$, and $y \in \mathcal{R}(x), y' \in \mathcal{R}(x')$, $\langle y' - y, x' - x \rangle \geq 0$. $\mathcal{R}$ is **pseudomonotone** iff for all $x, x' \in \mathcal{X}$, and $y \in \mathcal{R}(x), y' \in \mathcal{R}(x')$, $\langle y', x' - x \rangle \geq 0 \implies \langle y, x - x' \rangle \geq 0$. $\mathcal{R}$ is **quasimonotone** iff for all $x, x' \in \mathcal{X}$, and $y \in \mathcal{R}(x), y' \in \mathcal{R}(x')$, $\langle y', x' - x \rangle > 0 \implies \langle y, x - x' \rangle \geq 0$. We note the following relationship between these notions of monotonicity: $\text{monotone} \implies \text{pseudomonotone} \implies \text{quasimonotone}$.

2.2 Set Theory

2.2.1 Sets

The notion of a **set** (or collection or family) which are objects that consists of **elements** (or **points**) is taken as a primitive throughout, and as such a background in set theory is prerequisite for understanding the notions developed in this thesis. We refer the reader to the prologue of Folland (1999) for the necessary background.

We use the shorthands $\forall$ and $\exists$ to respectively mean for all (or for every), and there exists (or for some).

Unless otherwise noted, letters will be used as variables, i.e., placeholders that denote a value. We use caligraphic uppercase letters or Greek uppercase letters to denote sets (e.g., $\mathcal{X}$ or Φ).

$\mathcal{X} = \{u, v, w\}$ denotes the elements of the set $\mathcal{X}$, namely u, v, w. Order and repetitions are insignificant, that is, $\{u, v, w\} = \{u, w, w, v, u, w, v\}$.

The elipsis $\ldots$ is meant to be understood as a logical completion of any sequence, e.g., $\mathcal{X} = \{u, v, w, \ldots\} = \{u, v, w, x, y, z\}$.

$\emptyset \doteq \{\}$ denotes the empty set $\{\}$, i.e., the set without any elements.

$\mathcal{X}$ is called a **singleton** iff it contains only one element.

$[n]$ denotes the set of integers $\{1, \ldots, n\}$.

$[n : m]$ denotes the set of integers $\{n, n+1 \ldots, m\}$.

$\mathbb{N}$ denotes the set of natural numbers $\{1, 2, 3, \ldots\}$.

$\mathbb{R}$ denotes the set of real numbers (i.e., the set of all numbers strictly between $-\infty$ and ∞).

$\bar{\mathbb{R}}$ denotes the set of extended real numbers $\mathbb{R} \cup \{-\infty, \infty\}$.

We denote the positive and strictly positive (respectively, negative and strictly negative) elements of a set by $+$ and $++$ (respectively, $-$ and $--$) subscripts, respectively, e.g., $\mathbb{R}_+$ and $\mathbb{R}_{--}$.

$x \in \mathcal{X}$ means x is an element of the set $\mathcal{X}$.

$x \notin \mathcal{X}$ means x is not an element of the set $\mathcal{X}$.

$\mathcal{Y} \subseteq \mathcal{X}$ (or $\mathcal{X} \supseteq \mathcal{Y}$) means every element of $\mathcal{Y}$ is also an element of the set of $\mathcal{X}$, in which case we will say that $\mathcal{Y}$ is a subset of $\mathcal{X}$.

$\mathcal{Y} = \mathcal{X}$ means that $\mathcal{Y} \subseteq \mathcal{X}$ and $\mathcal{Y} \supseteq \mathcal{X}$.

$\mathcal{Y} \subset \mathcal{X}$ ($\mathcal{Y} \subsetneq \mathcal{X}$ or $\mathcal{X} \supset \mathcal{Y}$ or $\mathcal{X} \supsetneq \mathcal{Y}$) means that we have $\mathcal{Y} \subseteq \mathcal{X}$ but not $\mathcal{Y} = \mathcal{X}$, in which case we will say that $\mathcal{Y}$ is a strict subset of $\mathcal{X}$.

A set $\mathcal{X}$ is **non-empty** iff $\mathcal{X} \neq \emptyset$.

A set $\mathcal{X}$ is said to be **affine** iff for all $\lambda \in \mathbb{R}$ and $x, y \in \mathcal{X}$, we have $\lambda x + (1 - \lambda) y \in \mathcal{X}$.

A set $\mathcal{X}$ is said to be **convex** iff for all $\lambda \in [0, 1]$ and $x, y \in \mathcal{X}$, we have $\lambda x + (1 - \lambda) y \in \mathcal{X}$.

The set of all elements, which will be clear from context, is called the **universal set**, and is denoted by $\mathcal{U}$.

Given any set $\mathcal{X} \subseteq \mathcal{U}$, we denote its **power set**, i.e., the collection of all of its subsets, by $2^{\mathcal{X}} \doteq \{\mathcal{Y} \subseteq \mathcal{X}\}$.

$\{x \in \mathcal{X} \mid P(x)\}$ denotes the set of all elements $x \in \mathcal{X}$ for which the proposition $P(x)$ is true.

$\mathcal{X} \cup \mathcal{Y}$ is called the **union** of the sets $\mathcal{X}$ and $\mathcal{Y}$, and is given by the set $\{x \in \mathcal{U} \mid x \in \mathcal{X}$ or $x \in \mathcal{Y}\}$.

$\mathcal{X} \cap \mathcal{Y}$ is called the **intersection** of the sets $\mathcal{X}$ and $\mathcal{Y}$, and is given by the set $\{x \in \mathcal{U} \mid x \in \mathcal{X}$ and $x \in \mathcal{Y}\}$, often written as $\{x \mid x \in \mathcal{X}, x \in \mathcal{Y}\}$.

Two sets $\mathcal{X}, \mathcal{Y} \subseteq \mathcal{U}$ are said to be **disjoint** iff their intersection is empty, i.e., $\mathcal{X} \cap \mathcal{Y} = \emptyset$. A collection of sets $\mathcal{E} \subseteq \mathcal{U}$ is said to be **pairwise disjoint** iff for all $\mathcal{X}, \mathcal{Y} \in \mathcal{E}$ s.t. $\mathcal{X} \neq \mathcal{Y}$, $\mathcal{X}$ and $\mathcal{Y}$ are disjoint.

If $\mathcal{E} \subseteq \mathcal{U}$ is a collection of sets, we define the the union and intersection of its members respectively as:

$$\bigcup_{\mathcal{X} \in \mathcal{E}} \mathcal{X} \doteq \{x \in \mathcal{U} \mid \text{ for some } x \in \mathcal{X}\}$$
$$\bigcap_{\mathcal{X} \in \mathcal{E}} \mathcal{X} \doteq \{x \in \mathcal{U} \mid \text{ for all } x \in \mathcal{X}\}$$

$\{\mathcal{X}_i\}_{i \in \mathcal{I}}$ denotes a collection of $\mathcal{X}_i$'s, i.e., $\bigcup_{i \in \mathcal{I}} \{\mathcal{X}_i\}$. Similarly, for any $k, n \in \mathbb{N}$ s.t. $k \geq n$, $\{\mathcal{X}_i\}_{i=k}^n$ denotes the collection of the sets $\mathcal{X}_k, \ldots, \mathcal{X}_n$, i.e., $\bigcup_{i=k}^n \{\mathcal{X}_i\}$. When the index set $\mathcal{I}$ is clear from context, we will often write $\{\mathcal{X}_i\}_i$.

For any set $\mathcal{X} \subseteq \mathcal{M}$, and family of sets $\{\mathcal{Y}_i\}_{i \in \mathcal{I}}$ such that $\mathcal{X} \subseteq \bigcup_{i \in \mathcal{I}} \mathcal{Y}_i$, $\{\mathcal{Y}_i\}_{i \in \mathcal{I}}$ is called a **cover** of $\mathcal{X}$, and $\mathcal{X}$ is said to be **covered by** the $\mathcal{Y}_i$'s.

$\mathcal{X} \backslash \mathcal{Y}$ is called the **difference** of the sets $\mathcal{X}$ and $\mathcal{Y}$, and is given by $\{x \in \mathcal{U} \mid x \in \mathcal{X} \text{ and } x \notin \mathcal{Y}\}$.

The **complement** $\mathcal{X}^c$ of a set $\mathcal{X} \subseteq \mathcal{U}$ is given by the difference of the universal set $\mathcal{U}$ and $\mathcal{X}$, i.e., $\mathcal{X}^c \doteq \mathcal{U} \backslash \mathcal{X}$.

An n**-tuple** (or tuple when clear from context) $\boldsymbol{x} \doteq (x_1, x_2, \ldots, x_n)$ is an ordered array of $n \in \mathbb{N}_{++}$ elements s.t. two tuples $\boldsymbol{x}$ and $\boldsymbol{y}$ are equal, i.e., $(x_1, x_2, \ldots, x_n) = (y_1, y_2, \ldots, y_n)$, iff $x_i = y_i$ for all $i \in [n]$. We sometimes use a bold lowercase letter to stress that a variable is a tuple, but a normal font lowercase letter can also represent a tuple. $(x_i, \boldsymbol{x}_{-i})$ denotes the tuple $\boldsymbol{x} \doteq (x_1, \ldots, x_n)$, where $\boldsymbol{x}_{-i}$ denotes $\boldsymbol{x}$ with the ith element x_i removed. A 2-tuple is often also called a **pair**. For any $k, n \in \mathbb{N}$ s.t. $k \geq n$, $\{x_i\}_{i=k}^n$ denotes the tuple $(x_k, \ldots, x_n)$. When clear from context, we will often write $(x_i)_i$.

For convenience, we will denote an n-tuple which consists of the same number by the number in bold font with a subscript of n, e.g., $\mathbf{0}_3 \doteq (0, 0, 0)$ or $\mathbf{1}_4 \doteq (1, 1, 1, 1)$. When clear from, context, we will often omit the subscript.

We denote the n-dimensional ith **basis vector** $\boldsymbol{j}_i \doteq (1, \boldsymbol{0}_{n,-i})$.

$\mathcal{X} \times \mathcal{Y}$ is the **Cartesian product** of sets $\mathcal{X}$ and $\mathcal{Y}$ and is given by $\{(x, y) \mid x \in \mathcal{X}, y \in \mathcal{Y}\}$. The Cartesian product of a collection $\{\mathcal{X}_i\}_{i\in[n]}$ of $n \in \mathbb{N}$ sets is given by $\bigtimes_{i\in[n]} \mathcal{X}_i \doteq \mathcal{X}_1 \times \ldots \times \mathcal{X}_n$. When for all $i \in [n]$, $\mathcal{X}_i = \mathcal{Y}$, we write $\mathcal{Y}^n \doteq \bigtimes_{i\in[n]} \mathcal{X}_i$.

Δ_n denotes the unit simplex in $\mathbb{R}^n$, i.e., $\Delta_n = \{\boldsymbol{x} \in \mathbb{R}^n_+ \mid \sum_{i=1}^n x_i = 1\}$.

We denote the **affine hull** of any set $\mathcal{X}$ by $\operatorname{aff}(\mathcal{X}) \doteq \bigcap_{\mathcal{Y}\supset\mathcal{X}:\mathcal{Y} \text{ is affine}} \mathcal{Y}$. For $\mathcal{X} \subseteq \mathbb{R}^n$, by Caratheodory's theorem, this definition reduces to $\operatorname{aff}(\mathcal{X}) \doteq \{\sum_{i=1}^{n+1} \lambda_i x_i \mid \forall \boldsymbol{\lambda} \in \mathbb{R}^{n+1} \text{ s.t. } \sum_{i\in[n+1]} \lambda_i = 1\}$. We denote the **convex hull** of any set $\mathcal{X}$ by $\operatorname{conv}(\mathcal{X}) \doteq \bigcap_{\mathcal{Y}\supset\mathcal{X}:\mathcal{Y} \text{ is convex}} \mathcal{Y}$. For $\mathcal{X} \subseteq \mathbb{R}^n$, by Caratheodory's theorem, this definition reduces to $\operatorname{conv}(\mathcal{X}) \doteq \{\sum_{i=1}^{n+1} \lambda_i x_i \mid \forall \boldsymbol{\lambda} \in \Delta_{n+1}\}$.

For any real number $a \in \mathbb{R}$, $a\mathcal{X}$ denotes the (Minkowski) product, i.e., $a\mathcal{Y} \doteq \{ax \mid x \in \mathcal{X}\}$.

$\mathcal{X} + \mathcal{Y}$ denotes the (Minkowski) sum of $\mathcal{X}$ and $\mathcal{Y}$, i.e., $\mathcal{X} + \mathcal{Y} \doteq \{x + y \mid x \in \mathcal{X}, y \in \mathcal{Y}\}$.

$\mathcal{X} - \mathcal{Y}$ denotes the (Minkowski) difference of $\mathcal{X}$ and $\mathcal{Y}$, i.e., $\mathcal{X} - \mathcal{Y} \doteq \{x - y \mid x \in \mathcal{X}, y \in \mathcal{Y}\}$.

2.2.2 Relations

Relations Given sets $\mathcal{X}$ and $\mathcal{Y}$, a **relation** $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{Y}$ from $\mathcal{X}$ to its **codomain** $\mathcal{Y}$ is a subset of $\mathcal{X} \times \mathcal{Y}$. For any $x \in \mathcal{X}, y \in \mathcal{Y}$, we write $x \succeq_{\mathcal{R}} y$ to mean $(x, y) \in \mathcal{R}$.

The **domain** $\operatorname{dom}(\mathcal{R})$ of a relation $\mathcal{R}$ is given by the set $\operatorname{dom}(\mathcal{R}) \doteq \{x \in \mathcal{X} \mid \exists y \in \mathcal{Y} \text{ s.t. } x \succeq_{\mathcal{R}} y\}$.

The **range** (or **image**) of a relation $\mathcal{R}$ is given by the set $\operatorname{range}(\mathcal{R}) \doteq \{y \in \mathcal{Y} \mid \exists x \in \mathcal{X} \text{ s.t. } x \succeq_{\mathcal{R}} y\}$.

The **inverse** $\mathcal{R}^{-1} \subseteq \mathcal{Y} \times \mathcal{X}$ of a relation $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{Y}$ is $\mathcal{R}^{-1} \doteq \{(y, x) \in \mathcal{Y} \times \mathcal{X} \mid x \succeq_{\mathcal{R}} y\}$.

The **image** $\mathcal{R}(x) \subseteq \mathcal{Y}$ of $x \in \mathcal{X}$ under a relation $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{Y}$ is $\mathcal{R}(x) \doteq \{y \in \mathcal{Y} \mid x \succeq_{\mathcal{R}} y\}$.

The **image** $\mathcal{R}(\mathcal{A}) \subseteq \mathcal{Y}$ of a set $\mathcal{A} \subseteq \mathcal{X}$ under a relation $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{Y}$ is $\mathcal{R}(\mathcal{A}) \doteq \bigcup_{x\in\mathcal{A}} \mathcal{R}(x)$.

Let $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{Y}$ be a relation from $\mathcal{X}$ to $\mathcal{Y}$, and $\mathcal{R}'\mathcal{Y} \times \mathcal{Z}$ be a relation from $\mathcal{Y}$ to $\mathcal{Z}$. The **composition** $\mathcal{R}' \circ \mathcal{R}$ of $\mathcal{R}'$ with $\mathcal{R}$ is defined as:

$$\mathcal{R}' \circ \mathcal{R} \doteq \{(x, z) \in \mathcal{X} \times \mathcal{Z} \mid \exists y \in \mathcal{Y} \text{ s.t. } x \succeq_{\mathcal{R}} y, y \succeq_{\mathcal{R}'} z\}$$

If $\mathcal{Y} \doteq \mathcal{X}$, then $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{X}$ is said to be a **(binary) relation** on $\mathcal{X}$, in which case for all $x, y \in \mathcal{X}$, we say that x **succeeds** y iff $x \succeq_{\mathcal{R}} y$. Further, we say that x **preceeds** y and we write $x \preceq_{\mathcal{R}} y$ to mean $(y, x) \in \mathcal{R}$. We say that x is **similar** to y and write $x \simeq y$ iff $x \succeq_{\mathcal{R}} y$ and $y \preceq_{\mathcal{R}} x$, and we write $x \not\simeq y$ otherwise. We say that x **strictly succeeds** (respectively, **strictly preceeds**) y and write $x \succ_{\mathcal{R}} y$ (respectively, $x \prec_{\mathcal{R}} y$) to mean $x \succeq_{\mathcal{R}} y$ and $x \not\simeq_{\mathcal{R}} y$ (respectively, $x \preceq_{\mathcal{R}} y$ and $x \not\simeq_{\mathcal{R}} y$). When $\mathcal{R}$ is clear from context, we denote $\succeq_{\mathcal{R}}$, $\preceq_{\mathcal{R}}$, and $\simeq_{\mathcal{R}}$ simply by $\succeq$, $\preceq$, and $\simeq$, respectively.

For any binary relation $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{X}$, we also define the following properties. The relation $\mathcal{R}$ is **complete** iff for all $x, y \in \mathcal{X}$, either $x \succeq y$, $x \preceq y$, or $x \simeq y$. The relation $\mathcal{R}$ is **transitive** iff, for all $x, y, z \in \mathcal{X}$, $x \succeq z$ whenever $x \succeq y$ and $y \succeq z$. The relation $\mathcal{R}$ is **antisymmetric** iff for all $x, y \in \mathcal{X}$, $x \succeq y$ and $x \preceq y$, then $x = y$. The relation $\mathcal{R}$ is **reflexive** iff for all $x \in \mathcal{X}$, $x \succeq x$.

A **partial order** $(\mathcal{U}, \mathcal{R})$ consists of a universal set $\mathcal{U}$ and a binary relation $\mathcal{R} \subseteq \mathcal{U} \times \mathcal{U}$ which is transitive, antisymmetric, and reflexive. If, in addition, $\mathcal{R}$ is complete, then $(\mathcal{U}, \mathcal{R})$ is complete order. When $\mathcal{U} \doteq \mathbb{R}$, then we will assume that $(\mathcal{U}, \mathcal{R})$ is the usual (total) order on $\mathbb{R}$, in which case we denote $\succeq_{\mathcal{R}}, \preceq_{\mathcal{R}}, \simeq_{\mathcal{R}}, \succ_{\mathcal{R}}, \prec_{\mathcal{R}}$ by $\geq, \leq, =, >, <$. When $\mathcal{U} \doteq \mathbb{R}^n$, then we will assume the partial order $(\mathcal{U}, \mathcal{R})$ defined by the relation $\mathcal{R} \doteq \{(\boldsymbol{x}, \boldsymbol{y}) \in \mathbb{R}^n \times \mathbb{R}^n \mid x_i \geq y_i, \forall i \in [n]\}$, in which case, overloading notation, we denote $\succeq_{\mathcal{R}}, \preceq_{\mathcal{R}}, \simeq_{\mathcal{R}}, \succ_{\mathcal{R}}, \prec_{\mathcal{R}}$ by $\geq, \leq, =, >, <$.

For any partially ordered set $(\mathcal{U}, \mathcal{R})$, we define the **infimum** or **lower bound** (respectively, **supremum** or **upper bound**) of a set $\mathcal{X}$ as an element $x^* \in \mathcal{U}$ s.t. for all $x \in \mathcal{X}$, $x \succeq x^*$ (respectively, $x \preceq x^*$), in which case we denote $\inf(\mathcal{X}) \doteq x^*$ (respectively, $\sup(\mathcal{X}) \doteq x^*$). If the infimum $\inf(\mathcal{X})$ (respectively, supremum $\sup(\mathcal{X})$) is an element of $\mathcal{X}$, we then write $\min(\mathcal{X})$ (respectively, $\max(\mathcal{X})$).

2.2.3 Correspondences

A relation $\mathcal{R} \subset \mathcal{X} \times \mathcal{Y}$ from $\mathcal{X}$ to $\mathcal{Y}$ s.t. $\mathrm{dom}(\mathcal{R}) = \mathcal{X}$ is called a **correspondence** from $\mathcal{X}$ to its **codomain** $\mathcal{Y}$ and is denoted by $\mathcal{R} : \mathcal{X} \rightrightarrows \mathcal{Y}$. A **correspondence** $\mathcal{R} : \mathcal{X} \rightrightarrows \mathcal{Y}$ can be understood as a map[1] from $\mathcal{X}$ to subsets of $\mathcal{Y}$, in which for all $x \in \mathcal{X}$ we write $x \mapsto\mathrel{\mkern-14mu}\rightrightarrows \mathcal{R}(x) \subseteq \mathcal{Y}$. As a correspondence is a relation (i.e., a set), throughout this thesis, we will denote correspondences by calligraphic uppercase letters (e.g., $\mathcal{R}$). A correspondence is said to be **non-empty-valued** (respectively, **convex-valued**) iff for all $x \in \mathcal{X}$, $\mathcal{R}(x)$ is **non-empty** (respectively, **convex**).

2.2.4 Functions

A relation $f \subseteq \mathcal{X} \times \mathcal{Y}$ from $\mathcal{X}$ to $\mathcal{Y}$ with the property that for every $x \in \mathcal{X}$, there is a unique $y \in \mathcal{Y}$ s.t. $x \succeq_{\mathcal{R}} y$ is called a **function** (or **mapping**) from $\mathcal{X}$ to its **codomain** $\mathcal{Y}$ and is denoted by $f : \mathcal{X} \to \mathcal{Y}$. We call the element of y the value of f at x, and by abuse of notation, we define $x \mapsto f(x) \doteq y$ Thus, while the value of f at x (i.e., y), and its image at x (i.e., $\{y\}$) are both denoted by $f(x)$, the meaning of $f(x)$ will be clear from context. While a function $f : \mathcal{X} \to \mathcal{Y}$ should be understood as a relation (i.e., a set), we will often be working with the value (i.e., an element rather than a set) of f rather than its image $f(x)$, and as such we will often denote functions by a lowercase letter. In some cases, if we want to stress the type of the value of the function, we will denote the function by the type of its value. For instance, if the function is denoted by a bold lowercase letter (e.g., $\boldsymbol{f}$), then the mapping is vector valued, and if the function is denoted by bold uppercase letter (e.g., $\boldsymbol{F}$), then the mapping is matrix valued.

A function $f : \mathcal{X} \to \mathcal{Y}$ is said to be **affine** (or **linear**[2]) if for all $\alpha, \beta \in \mathbb{R}$ and $x, x' \in \mathcal{X}$ we have $f(\alpha x + \beta x') = \alpha f(x) + \beta f(x')$. We denote by $\mathbb{1}_{\mathcal{X}}(\boldsymbol{x})$ the indicator function (or Dirac

[1]When understood as a point-to-set (or multivalued) map, a **correspondence** $\mathcal{R} : \mathcal{X} \rightrightarrows \mathcal{Y}$ is denoted $\mathcal{R} : \mathcal{X} \to 2^{\mathcal{Y}}$ s.t. for all $x \in \mathcal{X}$, $\mathcal{R}$ maps x to subsets $\mathcal{Y}$, i.e., $x \mapsto\mathrel{\mkern-14mu}\rightrightarrows \mathcal{R}(x) \subseteq \mathcal{Y}$. While some authors have with authority argued the ill-posedness of such a definition (see page 1 of Dieudonné (1960)), such a view can be helpful in obtaining many theoretical results.

[2]We will not be making a distinction between the two concepts.

delta measure) for a set $\mathcal{X}$, with value 1 if $\boldsymbol{x} \in \mathcal{X}$ and 0 otherwise. We denote by $\chi_{\mathcal{X}}(\boldsymbol{x})$ the **characteristic function** for a set $\mathcal{X}$, with value 0 if $\boldsymbol{x} \in \mathcal{X}$ and ∞ otherwise.

2.2.5 Cardinality

Consider a function $f : \mathcal{X} \to \mathcal{Y}$ from $\mathcal{X}$ to $\mathcal{Y}$.

f is called **injective** (or **one-to-one**) iff $f(x) = f(\boldsymbol{x})$ implies $x = y$.

f is called **surjective** (or **onto**) iff $\mathrm{range}(f) = \mathcal{Y}$.

f is called **bijective** iff it is injective and surjective.

If $\mathcal{X}$ and $\mathcal{Y}$ are non-empty sets, we define the following expressions:

$$\mathrm{card}(\mathcal{X}) \leq \mathrm{card}(\mathcal{Y}) \quad \mathrm{card}(\mathcal{X}) \geq \mathrm{card}(\mathcal{Y}) \quad \mathrm{card}(\mathcal{Y}) = \mathrm{card}(\mathcal{X})$$

to mean that there exists $f : \mathcal{X} \to \mathcal{Y}$ which is injective, surjective or bijective respectively.

Additionally, we define $\mathrm{card}(\mathcal{X}) < \mathrm{card}(\mathcal{Y})$ (or $\mathrm{card}(\mathcal{Y}) > \mathrm{card}(\mathcal{X})$) to mean $\mathrm{card}(\mathcal{X}) \leq \mathrm{card}(\mathcal{Y})$ but not $\mathrm{card}(\mathcal{X}) = \mathrm{card}(\mathcal{Y})$.

Consider a set $\mathcal{X}$. $\mathcal{X}$ is **countable** iff $\mathrm{card}(\mathcal{X}) \leq \mathrm{card}(\mathbb{N})$ and **infinite** iff $\mathrm{card}(\mathcal{X}) \geq \mathrm{card}(\mathbb{N})$. $\mathcal{X}$ is **countably infinite** iff it is countable and infinite. A set $\mathcal{X}$ is said to be **finite** iff $\mathrm{card}(\mathcal{X}) < \mathrm{card}(\mathbb{N})$. A finite set $\mathcal{X}$ is said to have cardinality n iff $\mathrm{card}(\mathcal{X}) = \mathrm{card}([n])$, in which case we write $\mathrm{card}(\mathcal{X}) \doteq n$.[3]

2.2.6 Sequences

A function $f : \mathbb{N}_{++} \to \mathcal{X}$ from the set of positive integers $\mathbb{N}_{++}$ to a set $\mathcal{X}$ is called a **sequence** of points in $\mathcal{X}$, and denoted $\{x^{(n)}\}_n \doteq \{f(n)\}_n$ or $(x^{(n)})_n \doteq (f(n))_n$. While we denote a sequence as a tuple or a collection, it should always be understood as a function. Although we often denote elements of a sequence by superscripts with round brackets, e.g, $x^{(n)}$, to

[3] While this notation might suggest that $\mathrm{card}(\cdot)$ is a function, such an interpretation is only appropriate for finite sets for which cardinality can be interpreted as the number of elements in the the set, and should be avoided for infinite sets.

stress the sequential aspect of sequences, a sequence can also be denoted by lower or upper scripts without brackets, e.g., x_n and x^n.

2.2.7 Infinite Cartesian Product

Given an infinite collection of sets $\{\mathcal{X}_i\}_{i\in\mathcal{I}}$ (i.e., $\mathcal{I}$ is infinite), the **Cartesian product** of the $\mathcal{X}_i$s is defined as:

$$\bigtimes_{i\in\mathcal{I}} \mathcal{X}_i \doteq \left\{ f : \mathcal{I} \to \bigcup_{i\in\mathcal{I}} \mathcal{X}_i \right\} \tag{2.1}$$

We note that this definition of the Cartesian product for infinite collections of sets does not agree with the definition of the Cartesian product for finite collections of sets. As such, while for simplicity we will use the same notational convention for both, the definition of the operator $\bigtimes_{i\in\mathcal{I}}$ for infinite $\mathcal{I}$ should be understood as distinct from the one for finite $\mathcal{I}$. This distinction will be clear from context throughout this thesis. With this definition in hand, when for all $i \in \mathcal{I}$, $\mathcal{X}_i = \mathcal{Y}$, notice that $\bigtimes_{i\in\mathcal{I}} \mathcal{X}_i = \bigtimes_{i\in\mathcal{I}} \mathcal{Y} = \mathcal{Y}^{\mathcal{X}} = \{f : \mathcal{X} \to \mathcal{Y}\}$. As such, we denote $\mathcal{Y}^{\mathcal{X}} \doteq \{f : \mathcal{X} \to \mathcal{Y}\}$ the set of all functions from $\mathcal{X}$ to $\mathcal{Y}$.

2.3 Metric Spaces

A **metric space** is a tuple $(\mathcal{M}, d)$ that consists of a set $\mathcal{M}$, and a function $d : \mathcal{M} \times \mathcal{M} \to \mathbb{R}_+$ called a **metric** which takes as input any two points in $\mathcal{M}$ and outputs a value called the **distance** between those two points such that the following hold:

1. (Non-Degeneracy) $d(x, y) = 0$ iff $x = y$, $\forall x, y \in \mathcal{M}$
2. (Triangle Inequality) $d(x, z) \leq d(x, y) + d(y, z)$, $\forall x, y, z \in \mathcal{M}$
3. (Symmetry) $d(x, y) = d(y, x)$, $\forall x, y \in \mathcal{M}$

A sequence $\{x^{(n)}\}_{n\in\mathbb{N}} \subseteq \mathcal{M}$ is said to **converge** to some $x \in \mathcal{M}$ if for every $\varepsilon > 0$, there exists an integer $\overline{n} \in \mathbb{N}$ s.t. for all integers $m \geq \overline{n}$ we have that $d(x^{(m)}, x) \leq \varepsilon$. The point to which the sequence converges is called its **limit**. When the metric space is clear from context, if the sequence $\{x^{(n)}\}_{n\in\mathbb{N}} \subseteq \mathcal{M}$ converges to $x \in \mathcal{M}$, we then write $\lim_{n\to\infty} x^{(n)} = x$ or $x^{(n)} \to x$. A sequence $\{x^{(n)}\}_{n\in\mathbb{N}} \subset \mathcal{M}$ is said to be a **Cauchy sequence** if for all $\epsilon > 0$, there exists $\overline{n} \in \mathbb{N}$ such that for all integers $n, m > \overline{n}$, we have $d(x^{(n)}, x^{(m)}) < \epsilon$.

A set $\mathcal{X}$ is said to be **closed** if for any convergent sequence $\{x^{(n)}\}_{n\in\mathbb{N}} \subseteq \mathcal{X}$ s.t. $x^{(n)} \to x \in \mathcal{M}$, we have $x \in \mathcal{X}$. A set $\mathcal{X}$ is said to be **open** if its complement $\mathcal{M} \setminus \mathcal{X}$ is closed. For any set $\mathcal{X}$, we define the distance of any point x to the set $\mathcal{X}$ as $d(x, \mathcal{X}) \doteq \min_{x'\in\mathcal{X}} d(x, x')$. We also define the diameter of a set by $\mathrm{diam}(\mathcal{X}) \doteq \max_{x,x'\in\mathcal{X}} d(x, x')$. A set $\mathcal{X}$ is said to be **bounded** iff $\mathrm{diam}(\mathcal{X}) < \infty$.

Throughout this thesis, we will be concerned only with complete metric spaces. A subset $\mathcal{X} \subseteq \mathcal{M}$ of $\mathcal{M}$ is said to be **complete** if every Cauchy sequence $\{x^{(n)}\}_{n\in\mathbb{N}} \subset \mathcal{X}$ converges, i.e., $x^{(n)} \to x \in \mathcal{M}$, and its limit $x \in \mathcal{X}$. A metric space $(\mathcal{M}, d)$ is said to be **complete** if $\mathcal{M}$ is complete. We note that any closed subset of a complete metric space is complete. A common example of a complete metric space is the **Euclidean (metric) space** $(\mathbb{R}^n, d)$ with the **Euclidean metric** $d(\boldsymbol{x}, \boldsymbol{y}) \doteq \sqrt{\sum_{i=1}^n (x_i - y_i)^2}$.

For any $\varepsilon \geq 0$, we write $\mathcal{B}_\varepsilon[x] = \{x' \in \mathcal{M} \mid d(x, x') \leq \varepsilon\}$ and $\mathcal{B}_\varepsilon(x) = \{x' \in \mathcal{M} \mid d(x, x') < \varepsilon\}$ to denote the closed and open ε-ball centered at $x \in \mathcal{M}$, respectively. A point $x \in \mathcal{X}$ is called an interior point of $\mathcal{X}$, if there exists $\varepsilon > 0$ s.t. $\mathcal{B}_\varepsilon(x) \subseteq \mathcal{X}$. The **interior** $\operatorname{int}(\mathcal{X})$ of a set $\mathcal{X}$ consists of the set of all interior points of $\mathcal{X}$. The **relative interior** of a set $\mathcal{X}$ consists of the interior of $\mathcal{X}$ within its affine hull, i.e., $\operatorname{relint}(\mathcal{X}) \doteq \{x \in \mathcal{X} \mid \exists \varepsilon > 0, \mathcal{B}_\varepsilon(x) \cap \operatorname{aff}(\mathcal{X}) \subseteq \mathcal{X}\}$ A set $\mathcal{X} \subseteq \mathcal{M}$ is said to be **totally bounded** if for every $\varepsilon > 0$, it can be covered by finitely many open ε-balls. A set $\mathcal{X} \subseteq \mathcal{M}$ is said to be **compact** if it is complete and totally bounded. By the Heine-Borel Theorem, any closed and bounded subset of $\mathbb{R}^n$ is compact.

A relation $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{Y}$ is **upper** (or **outer**[4]) **hemicontinuous** if for any sequence $\{(x^{(n)}, y^{(n)})\}_{n \in \mathbb{N}_+} \subset \mathcal{X} \times \mathcal{Y}$ such that $(x^{(n)}, y^{(n)}) \to (x, y)$ and $x^{(n)} \succeq y^{(n)}$ for all $n \in \mathbb{N}_+$, it also holds that $x \succeq y$. Note that if $\mathcal{R}$ is a compact set, then it is upper hemicontinuous. A relation $\mathcal{R} \subseteq \mathcal{X} \times \mathcal{Y}$ is **inner** (or **lower**) **hemicontinuous** if for any $y \in \mathcal{Y}$ and sequence $\{x^{(n)}\}_{n \in \mathbb{N}_+} \subseteq \mathcal{X}$ such that $x^{(n)} \to x$ and $x \succeq y$, there exists $\{y^{(n)}\}_{n \in \mathbb{N}_+} \subseteq \mathcal{Y}$ s.t. $x^{(n)} \succeq y^{(n)}$ for all $n \in \mathbb{N}_+$ and $y^{(n)} \to y$. A relation is said to be **continuous** if it is both upper and lower hemicontinuous, or equivalently if for any sequence $\{(x^{(n)}, y^{(n)})\}_{n \in \mathbb{N}_+} \subset \mathcal{X} \times \mathcal{Y}$ s.t. $x^{(n)} \to x$ and $x^{(n)} \succeq y^{(n)}$ for all $n \in \mathbb{N}_+$, it also holds that $y^{(n)} \to y$ and $x \succeq y$.

Since any correspondence is a relation, we can define (upper and lower) hemicontinuity for correspondences analogously. A correspondence $\mathcal{R} : \mathcal{X} \rightrightarrows \mathcal{Y}$ is said to be **upper** (or **outer**) **hemicontinuous** if for any sequence $\{(x^{(n)}, y^{(n)}\}_{n \in \mathbb{N}_+} \subset \mathcal{X} \times \mathcal{Y}$ such that $(x^{(n)}, y^{(n)}) \to (x, y)$ and $y^{(n)} \in \mathcal{R}(x^{(n)})$ for all $n \in \mathbb{N}_+$, it also holds that $y \in \mathcal{R}(x)$. A correspondence $\mathcal{R} : \mathcal{X} \rightrightarrows \mathcal{Y}$ is said to be **lower** (or **inner**) **hemicontinuous** if for any $y \in \mathcal{Y}$ and sequence $\{x^{(n)}\}_{n \in \mathbb{N}_+} \subset \mathcal{X}$ such that $x^{(n)} \to x$ and $y \in \mathcal{R}(x)$, there exists $\{y^{(n)}\}_{n \in \mathbb{N}_+} \subset \mathcal{Y}$ s.t. $y^{(n)} \in \mathcal{R}(x^{(n)})$ for all $n \in \mathbb{N}_+$ and $y^{(n)} \to y$. A correspondence $\mathcal{R} : \mathcal{X} \rightrightarrows \mathcal{Y}$ is **continuous** if for any sequence $\{x^{(n)}\}_{n \in \mathbb{N}_+} \subset \mathcal{X}$ such that $x^{(n)} \to x$, it holds that $\mathcal{R}(x^{(n)}) \to \mathcal{R}(x)$. A correspondence is said to be **closed-valued** (respectively, **compact-valued** / **convex-valued** / **singleton-valued**) iff for all $x \in \mathcal{X}$, $\mathcal{R}(x)$ is closed (respectively, compact / convex / a singleton).

For functions, the analogous definitions of upper and lower hemicontinuous relations can be shown be equivalent, and as such considering only continuity becomes enough. In

[4]Note that certain authors make a distinction between upper and outer hemicontinuity by adopting a weaker definition of outer hemicontinuity (see, for instance Border (2010)). We will use these two terms interchangeably.

particular, a function f is **continuous** if for all sequences $\{x^{(n)}\}_{n\in\mathbb{N}}$ s.t. $x^{(n)} \to x \in \mathcal{X}$, it holds that $f(x^{(n)}) \to f(x)$. Equivalently, a function $f : \mathcal{X} \to \mathcal{Y}$ is said to be **continuous** if for all $x \in \mathcal{X}$ and $\varepsilon > 0$, there exists $\delta > 0$ s.t. for all $y \in \mathcal{X}$, if $d(x, y) < \delta$, then $d(f(x), f(y)) < \epsilon$.

2.4 Normed Spaces

A **normed (vector) space** is a tuple $(\mathcal{X}, \|\cdot\|)$ that consists of a **(vector) space** $\mathcal{X}$ and a function $\|\cdot\|: \mathcal{X} \to \mathbb{R}_+$ called a **norm** such that the following hold:

1. (Normalized) $\|x\|= 0$ iff $x = 0, \forall x \in \mathcal{X}$

2. (Homogeneity) $\|cx\|= |c|\|x\|, \forall c \in \mathbb{R}, x \in \mathcal{X}$

3. (Triangle Inequality) $\|x + y\|\leq \|x\|+||y||, \forall x, y \in \mathcal{X}$

Note that $\mathcal{X}$ can be any set that satisfies the axioms of vector spaces (e.g., $\mathcal{X}$ can be a set of functions). Any normed space $(\mathcal{X}, \|\cdot\|)$ is a metric space since any norm $\|\cdot\|$ defines the **induced metric** $d(x, y) = \|x - y\|$ such that $(\mathcal{X}, d)$ is a valid metric space called the **metric space induced by** $(\mathcal{X}, \|\cdot\|)$. That is, any normed space is also a metric space, and as such the definitions provided in Section 2.3 all apply to normed spaces.

The canonical example of a normed vector space is the ℓ_n^p normed space $(\mathbb{R}^n, \|\cdot\|_p)$ defined by the p-norm $\|\boldsymbol{x}\|_p\doteq \sqrt[p]{\sum_{i=1}^n x_i^p}$. For $p \to \infty$, we obtain the **uniform** (or **sup**) **norm** $\|\cdot\|_\infty$, which is defined as $\|\boldsymbol{x}\|_\infty\doteq \max_{i\in[n]}\{|x_i|\}$. Throughout this thesis, we will mostly be working with the **Euclidean (normed) space** (or ℓ_n^2 **normed space**) $(\mathbb{R}^n, \|\cdot\|_2)$. If the metric space $(\mathcal{X}, d)$ induced by $(\mathcal{X}, \|\cdot\|)$ is complete, then $(\mathcal{X}, \|\cdot\|)$ is called a **Banach space** (or a **complete normed space**). Note that any ℓ_n^p space is a Banach space.

Let $(\mathcal{X}, \|\cdot\|_\mathcal{X})$ and $(\mathcal{Y}, \|\cdot\|_\mathcal{Y})$ be **normed spaces**. A function $f : (\mathcal{X}, \|\cdot\|_\mathcal{X}) \to (\mathcal{Y}, \|\cdot\|_\mathcal{Y})$ or correspondence $\mathcal{R} : (\mathcal{X}, \|\cdot\|_\mathcal{X}) \rightrightarrows (\mathcal{Y}, \|\cdot\|_\mathcal{Y})$ between normed spaces is called an **operator**. We often write $f : \mathcal{X} \to \mathcal{Y}$ and $\mathcal{R} : \mathcal{X} \rightrightarrows \mathcal{Y}$, respectively, when the normed space is clear from context.

A seminal theorem in functional analysis is the Kakutani-Glicksberg Fixed Point Theorem,[5] which provides sufficient conditions for a **fixed point** of a correspondence $\mathcal{R} : \mathcal{X} \rightrightarrows \mathcal{X}$, i.e., a point $x \in \mathcal{X}$ s.t. $x \in \mathcal{R}(x)$ to exist.

Theorem 2.4.1 [Kakutani-Glicksberg Fixed Point Theorem (Kakutani, 1941; Glicksberg, 1952)]**.**
Consider a normed space $(\mathcal{U}, \|\cdot\|)$, and a non-empty, compact, convex set $\mathcal{X} \subseteq \mathcal{U}$. If $\mathcal{R} : \mathcal{X} \rightrightarrows \mathcal{X}$ is upper hemicontinuous, and non-empty-, compact, and convex-valued, then there exists a fixed point $x \in \mathcal{X}$ s.t. $x \in \mathcal{R}(x)$.

[5]While Kakutani was concerned only with Euclidean spaces (Kakutani, 1941), Glicksberg subsequently generalized Kakutani's results to convex Hausdorff linear topological spaces (see Section 1 of Glicksberg (1952) for the relevant definitions). Hence, since any normed (vector) space is a convex Hausdorff linear topological space (see, for instance Section 5 of Folland (1999)), we state Glicksberg's result for the special case of normed spaces, which suffices for all applications of this theorem in this thesis. Since any Euclidean space is a normed space, the version of the theorem stated here generalizes Kakutani's to non-Euclidean spaces (e.g., functional spaces). Additionally, note that Glicksberg (1952) states his result for closed correspondences, but as the domain and range of the correspondence is compact-valued, it suffices to assume upper hemicontinuity instead, as any upper hemicontinuous and compact-valued correspondence is closed (see Section 2.3).

2.5 Inner Product Spaces

An **inner product space** is a tuple $(\mathcal{X}, \langle\cdot,\cdot\rangle)$ that consists of a vector space $\mathcal{X}$ and a function $\langle\cdot,\cdot\rangle : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ called an **inner product** (or **scalar product**) such that the following hold:

1. (Normalized) $\langle x, x\rangle > 0, \forall x \neq 0$
2. (Symmetry) $\langle x, y\rangle = \langle y, x\rangle, \forall x, y \in \mathcal{X}$
3. (Bilinear) $\langle ax + bx', y\rangle = a\langle x, y\rangle + b\langle x', y\rangle, \forall a, b \in \mathbb{R} \in\ x, x', y \in \mathcal{X}$

Any inner product space $(\mathcal{X}, \langle\cdot,\cdot\rangle)$ is a normed space (and hence, in turn, a metric space) since any inner product $\langle\cdot,\cdot\rangle$ defines the **induced norm** $\|x\| \doteq \sqrt{\langle x, x\rangle}$ such that $(\mathcal{X}, \langle\cdot,\cdot\rangle)$ is a valid normed (vector) space called a **pre-Hilbert space** $(\mathcal{X}, \|\cdot\|)$. As such, the definitions provided in Section 2.3 and Section 2.4 all apply to inner product spaces. A pre-Hilbert space that is complete is called a **Hilbert space**. Throughout this thesis, unless otherwise mentioned, we will be working with the ℓ_n^2 inner product spaces $(\mathbb{R}^n, \langle\cdot,\cdot\rangle)$ where $\langle \boldsymbol{x}, \boldsymbol{y}\rangle \doteq \sum_{i=1}^n x_i y_i$, whose Hilbert space is given by the ℓ_n^2 normed space which itself corresponds to the Eucliean metric space.

2.6 Measure and Probability Spaces

A **measurable space** $(\mathcal{X}, \mathcal{F})$ consists of a set $\mathcal{X}$ and a collection $\mathcal{F}$ of subsets of $\mathcal{X}$ which satisfies the following conditions:

1. (Closure under finite unions) for all $\mathcal{B}_1, \ldots \mathcal{B}_n \in \mathcal{F}$, $\bigcup_{i=1}^n \mathcal{B}_i \in \mathcal{F}$

2. (Closure under complements) for all $\mathcal{B} \in \mathcal{F}$, $\mathcal{B}^c \in \mathcal{F}$,

Let $(\mathcal{X}, \mathcal{F}_\mathcal{X})$ and $(\mathcal{Y}, \mathcal{F}_\mathcal{Y})$ be two measurable spaces. A function $f : \mathcal{X} \to \mathcal{Y}$ is said to be a measurable iff for all $\mathcal{B} \in \mathcal{F}_\mathcal{Y}$, the inverse image of $\mathcal{B}$ is contained in $\mathcal{F}_\mathcal{X}$, i.e., $f^{-1}(\mathcal{B}) \in \mathcal{F}_\mathcal{X}$. If $f : \mathcal{X} \to \mathcal{Y}$ is a measurable function, we will often write $f : (\mathcal{X}, \mathcal{F}_\mathcal{X}) \to (\mathcal{Y}, \mathcal{F}_\mathcal{Y})$.

A **measure** $\mu : \mathcal{F} \to \mathbb{R}_+$ on a measurable space $(\mathcal{X}, \mathcal{F})$ is a function which satisfies the following conditions:

1. (Normalized) $\mu(\emptyset) = 0$

2. (Countable additivity) for all pairwise disjoint collection of sets $\{\mathcal{B}_i\}_{i=1}^\infty \subseteq \mathcal{F}$, $\mu(\bigcup_{i=1}^\infty \mathcal{B}_i) = \sum_{i=1}^\infty \mu(\mathcal{B}_i)$

A **measure space** $(\mathcal{X}, \mathcal{F}, \mu)$ is a triple which consists of a measurable space $(\mathcal{X}, \mathcal{F})$ and a measure μ on $(\mathcal{X}, \mathcal{F})$. Let $\mathcal{X}$ be any set and let $(\mathcal{X}, \mathcal{F})$ be an associated measurable space. We write $\Delta(\mathcal{X}, \mathcal{F}) \doteq \{\mu : (\mathcal{X}, \mathcal{F}) \to [0, 1]\}$ to denote the set of probability measures on $(\mathcal{X}, \mathcal{F})$. When $\mathcal{F}$ is clear from context, we simply write $\Delta(\mathcal{X})$. We also define the **support** of a measure $\mu \in \Delta(\mathcal{X})$ as $\mathrm{supp}(\mu) \doteq \{x \in \mathcal{X} : \mu(x) > 0\}$.

A **simple function** $s : \mathcal{X} \to \mathbb{R}_+$ is a measurable function of the form $s(x) \doteq \sum_{i=1}^n \alpha_i \mathbb{1}_{\mathcal{Y}_i}(x)$ for some $\{\alpha_i\}_i \subseteq \mathbb{R}_+$ and $\{\mathcal{Y}_i\}_i \subseteq \mathcal{F}$. The (Lebesgue) integral of a simple function s over a set $\mathcal{B} \in \mathcal{F}$ is defined as:

$$\int_{x \in \mathcal{B}} s(x) d\mu(x) \doteq \sum_{i=1}^n \alpha_i \mu(\mathcal{Y}_i \cap \mathcal{B}) \tag{2.2}$$

The (Lebesgue) integral of a positive measurable function $f : \mathcal{X} \to \mathbb{R}_+$ is defined as:

$$\int_{x \in \mathcal{B}} f(x) d\mu(x) \doteq \sup \left\{ \int_{x \in \mathcal{B}} s(x) d\mu(x) \mid s \text{ is a simple function and } \forall x \in \mathcal{X}, s(x) \leq f(x) \right\} \tag{2.3}$$

We then extend this definition of the (Lebesgue) integral to any measurable function $f : \mathcal{X} \to \mathbb{R}$ by defining:

$$\int_{x \in \mathcal{B}} f(x) d\mu(x) \doteq \int_{x \in \mathcal{B}} \max\{f(x), 0\} \, d\mu(x) - \int_{x \in \mathcal{B}} \max\{-f(x), 0\} \, d\mu(x) \tag{2.4}$$

When clear from context, we will often denote $\int_{x \in \mathcal{B}} f(x) d\mu(x)$ by $\int_{\mathcal{B}} s d\mu$ or $\int_{\mathcal{B}} f(x) dx$.

A measurable function $f : \mathcal{X} \to \mathbb{R}$ is **integrable** on $\mathcal{B}$ iff $\int_{x \in \mathcal{B}} |f(x)| d\mu(x) < \infty$, where $x \mapsto |x|$ is the absolute value function. We note that any bounded measurable function $f : \mathcal{X} \to \mathbb{R}$ is integrable on $\mathcal{B}$ if $\mu(\mathcal{B}) < \infty$. If f is integrable on $\mathcal{X}$ itself, then f is said to be **integrable**.

A **probability space** is a measure space $(\mathcal{O}, \mathcal{E}, \mu)$ where

1. $\mathcal{O}$ is called the **sample space** and its elements are called **outcomes**
2. $\mathcal{E}$ is the event space which consists of sets of outcomes
3. $\mu : \mathcal{F} \to [0, 1]$ is a probability measure which satisfies $\mu(\mathcal{F}) = 1$

Given a measurable space $(\mathcal{X}, \mathcal{F})$ and a probability space $(\mathcal{O}, \mathcal{E}, \mu)$, a **random variable** is a measurable function $X : (\mathcal{O}, \mathcal{E}) \to (\mathcal{X}, \mathcal{F})$. A random variable maps outcomes in a probability space to elements of a measurable space, allowing us to quantify the occurrence of random outcomes. Throughout this thesis, we will denote random variables with normal font capital letters.

The probability that a random variable X takes on a value $x \in \mathcal{X}$ is denoted by:

$$\mathbb{P}_{X \sim \mu}(X = x) \doteq \mu(\{o \in \mathcal{O} \mid X(o) = x\}) \tag{2.5}$$

Similarly, the probability that a random variable X takes on a value in the set $\mathcal{Y} \subseteq \mathcal{X}$ is denoted by:

$$\mathbb{P}_{X \sim \mu}(X \in \mathcal{Y}) \doteq \mu(\{o \in \mathcal{O} \mid X(o) \in \mathcal{Y}\}) \tag{2.6}$$

The **expectation** of a random variable X is defined as:

$$\mathop{\mathbb{E}}_{X \sim \mu}[X] \doteq \int_{\mathcal{O}} X d\mu \tag{2.7}$$

2.7 (Sub)differential Calculus

Let $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ and $(\mathcal{Y}, \|\cdot\|_{\mathcal{Y}})$ be **normed spaces**. Consider an operator $f : (\mathcal{X}, \|\cdot\|_{\mathcal{X}}) \to (\mathcal{Y}, \|\cdot\|_{\mathcal{Y}})$. A linear operator $f : \mathcal{X} \to \mathcal{Y}$ is said to be **bounded** if there exists $c < \infty$ s.t. for all $x \in \mathcal{X}$,

$$\|f(x)\|_{\mathcal{X}} \leq c\|x\|_{\mathcal{Y}} \tag{2.8}$$

The **dual (vector) space** $\mathcal{X}^*$ of any (vector) space $\mathcal{X}$ consists of all linear functions $f : \mathcal{X} \to \mathbb{R}$ and is associated with a **dual normed vector space** $(\mathcal{X}^*, \|\cdot\|_{\mathcal{X}}^*)$ where $\|f\|_{\mathcal{X}}^* \doteq \sup_{\substack{x \in \mathcal{X} \\ \|x\| \leq 1}} \|f(x)\|_{\mathcal{X}}$.

The **directional (or Gâteau) derivative** (Gâteaux, 1913) of an operator $f : \mathcal{X} \to \mathcal{Y}$ at $\widehat{x} \in \mathcal{X}$ in the direction of $a \in \mathcal{X}$ is a linear function $\nabla_{\widehat{x}} f(a) \in \mathcal{X}^*$ on $\mathcal{X}$ s.t.:

$$\nabla_a f(\widehat{x}) = \lim_{t \to 0} \frac{f(\widehat{x} + ta) - f(\widehat{x})}{t} \tag{2.9}$$

f is said to be **differentiable**, if for all $\widehat{x}, a \in \mathcal{X}$, the directional derivative $\nabla_{\boldsymbol{a}} f(\widehat{x})$ exists. f is said to be **continuously differentiable** if for all $a \in \mathcal{X}$, $x \mapsto \nabla_{\boldsymbol{a}} f(x)$ is continuous. If $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \doteq \mathbb{R}$, overloading notation, we define the **partial derivative** of the function $f : \mathcal{X} \to \mathbb{R}$ w.r.t. $x_i \in \mathbb{R}$ for all $i \in [n]$ at $\widehat{\boldsymbol{x}} \in \mathcal{X}$ as $\nabla_{x_i} f(\widehat{\boldsymbol{x}}) \doteq \nabla_{\boldsymbol{j}_i} f(\widehat{\boldsymbol{x}})$, and the gradient (or Fréchet derivative) of $f : \mathcal{X} \to \mathbb{R}$ at $\widehat{\boldsymbol{x}} \in \mathcal{X}$ as $\nabla f(\widehat{\boldsymbol{x}}) = \nabla_{\boldsymbol{x}} f(\widehat{\boldsymbol{x}}) \doteq (\nabla_{x_i} f(\widehat{\boldsymbol{x}}))_{i=1}^n$.

The **(Clarke) subdifferential** (Clarke, 1990) of a function $f : \mathcal{X} \to \mathcal{Y}$ is a set $\mathcal{D} f(x) \subseteq \mathcal{X}^*$ of linear functions on $\mathcal{X}$ defined as $\mathcal{D} f(x) \doteq \operatorname{conv}\left\{\lim_{k \to \infty} \nabla f(x^{(k)}) \mid \exists x^{(k)} \to x \text{ s.t. } x^{(k)} \in \operatorname{dom}(\nabla f)\right\}$. Analogously, we define the **directional (Clarke) subdifferential** $\mathcal{D}_x$ w.r.t. $x \in \mathcal{X}$ by replacing the gradient operator in the definition by the directional derivative, i.e., ∇_x. To simplify notation, we often write $\partial_x f(\widehat{x})$ to refer to an arbitrary **subgradient** (i.e., an element of the subdifferential) of f at x, e.g., $\partial_x f(\widehat{x}) \in \mathcal{D}_x f(\widehat{x})$.

When f is continuously differentiable, by the definition of continuity, the subdifferential is singleton-valued and for all $\widehat{x} \in \mathcal{X}$, $\mathcal{D}_x f(\widehat{x}) \doteq \{\nabla_x f(\widehat{x})\}$. A function f is said to be **subdifferentiable** iff its subdifferential is non-empty for all points in its domain, i.e., for all $x \in \mathcal{X}$, $\partial_x f(x) \neq \emptyset$. We note that a function is subdifferentiable if it is locally-Lipschitz continuous[6] (Clarke, 1990). Throughout this thesis we will work only with subdifferentiable functions; as such, for any function $f : \mathcal{X} \to \mathcal{Y}$, we define its **subdifferential correspondence** $\mathcal{D}f : \mathcal{X} \rightrightarrows \mathcal{Y}$ as the correspondence that takes as input a point in the domain $x \in \mathcal{X}$, and outputs the subdifferential $\mathcal{D}f(x)$ of f at x. We define the directional subdifferential correspondence $\mathcal{D}_x f : \mathcal{X} \rightrightarrows \mathcal{Y}$ of any subdifferentiable function f similarly. We also note that for any Lipschitz-continuous function f, the subdifferential correspondence $\mathcal{D}f$ is upper hemicontinuous, non-empty-, and compact-valued (Clarke, 2007). For any continuous and convex function $f : \mathcal{X} \to \mathbb{R}$, the subdifferential correspondence $\mathcal{D}f$ is upper hemicontinuous, non-empty-, compact-, and convex-valued (see Theorem 24.4 of Pryce (1973)).

[6]We refer the reader to Section 2.8 for a definition.

2.8 Primitive Function Structures

In this section, we introduce the key definitions that will be used to derive the results in this thesis. These definitions are refinements of the notions of continuity and convexity for functions, and monotonicity for correspondences.

Consider a complete order $(\mathcal{U}, \mathcal{R})$ and a function $f : \mathcal{U} \to \mathbb{R}$. f is said to be **increasing** (respectively, **decreasing**) over $\mathcal{X} \subseteq \mathcal{U}$ iff for all $x, y \in \mathcal{X}$ s.t. $x \succeq_{\mathcal{R}} y$, $f(x) \geq f(y)$ (respectively, $f(x) \leq f(y)$). f is said to be **strictly increasing** (respectively, **strictly decreasing**) iff for all $x, y \in \mathcal{X}$ s.t. $x \succ_{\mathcal{R}} y$, $f(x) > f(y)$ (respectively, $f(x) < f(y)$).

Monotonicity Properties of Correspondences

To obtain our computational results, we will rely on generalized monotonicity properties of correspondences. Let $(\mathcal{X}, \langle \cdot, \cdot \rangle)$ be an inner product space. Consider a correspondence $\mathcal{R} : \mathcal{X} \rightrightarrows \mathcal{X}$.

Definition 2.8.1 [Weakly-Monotone/Dissipative Correspondences]**.**
$\mathcal{R}$ is μ**-weakly-monotone** with **modulo of monotonocity** $\mu \in \mathbb{R}$ iff

$$\left\langle y' - y, x' - x \right\rangle \geq -\mu \|x - x'\|^2 \qquad\qquad \forall y \in \mathcal{R}(x), y' \in \mathcal{R}(x')$$

$\mathcal{R}$ is ν**-weakly-dissipative** with **modulo of dissipativity** $\nu \in \mathbb{R}$ iff

$$\left\langle y' - y, x' - x \right\rangle \leq \nu \|x - x'\|^2 \qquad\qquad \forall y \in \mathcal{R}(x), y' \in \mathcal{R}(x')$$

Note that in the above definition when $\mu < 0$ (respectively, $\nu < 0$), a μ-weakly-monotone (respectively, ν-weakly-dissipative) correspondence is often called $(-\mu)$**-strongly-monotone** (respectively, $(-\nu)$**-strongly-dissipative**.

In the special case that $\mu \doteq 0$, and $\nu \doteq 0$, we recover the definitions of monotone and dissipative operators:

Definition 2.8.2 [Monotone/Dissipative Correspondences]**.**

$\mathcal{R}$ is **monotone** iff:

$$\langle y' - y, x' - x \rangle \geq 0 \qquad\qquad \forall y \in \mathcal{R}(x), y' \in \mathcal{R}(x')$$

$\mathcal{R}$ is **dissipative** iff:

$$\langle y' - y, x' - x \rangle \leq 0 \qquad\qquad \forall y \in \mathcal{R}(x), y' \in \mathcal{R}(x')$$

Definition 2.8.3 [Pseudomonotone/Pseudodissipative]**.**

$\mathcal{R}$ is **pseudomonotone** iff:

$$\langle y', x' - x \rangle \geq 0 \text{ implies } \langle y, x - x' \rangle \geq 0 \qquad\qquad \forall y \in \mathcal{R}(x), y' \in \mathcal{R}(x')$$

$\mathcal{R}$ is **pseudodissipative** iff:

$$\langle y', x' - x \rangle \leq 0 \text{ implies } \langle y, x - x' \rangle \leq 0 \qquad\qquad \forall y \in \mathcal{R}(x), y' \in \mathcal{R}(x')$$

Definition 2.8.4 [quasimonotone/quasidissipative]**.**

Consider a correspondence $\mathcal{R} : \mathcal{X} \rightrightarrows \mathcal{X}$.

$\mathcal{R}$ is **quasimonotone** iff:

$$\langle y', x' - x \rangle > 0 \text{ implies } \langle y, x - x' \rangle \geq 0 \qquad\qquad \forall y \in \mathcal{R}(x), y' \in \mathcal{R}(x')$$

$\mathcal{R}$ is **quasidissipative** iff:

$$\langle y', x' - x \rangle < 0 \text{ implies } \langle y, x - x' \rangle \leq 0 \qquad\qquad \forall y \in \mathcal{R}(x), y' \in \mathcal{R}(x')$$

We note the following relationships among these properties:

$$\text{monotone} \implies \text{pseudomonotone} \implies \text{quasimonotone} \tag{2.10}$$

$$\text{dissipative} \implies \text{pseudodissipative} \implies \text{quasidissipative} \tag{2.11}$$

Lipschitz Properties of Functions

Let $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ and $(\mathcal{Y}, \|\cdot\|_{\mathcal{Y}})$ be normed spaces. Consider a function $f : (\mathcal{X}, \|\cdot\|_{\mathcal{X}}) \to (\mathcal{Y}, \|\cdot\|_{\mathcal{Y}})$. We will first consider refinements of the notions of continuity and continuous differentiability, which we will use to derive the complexity results in this thesis.

Definition 2.8.5 [Lipschitz-Continuity]**.**

A function $f : \mathcal{X} \to \mathcal{Y}$ is said to be ℓ_f-**Lipschitz-continuous on** $\mathcal{S} \subseteq \mathcal{X}$ iff for all $x_1, x_2 \in \mathcal{X}$

$$\|f(x_1) - f(x_2)\|_{\mathcal{Y}} \leq \ell_f \|x_1 - x_2\|_{\mathcal{X}} , \tag{2.12}$$

When $\mathcal{S} = \mathcal{X}$, then f is simply called ℓ_f-**Lipschitz-continuous**.

We note that any continuously differentiable function f on a non-empty and compact set $\mathcal{S}$ is guaranteed to be ℓ-Lipschitz-continuous on $\mathcal{S}$ with $\ell \doteq \max_{x \in \mathcal{S}} \|\nabla f(x)\|_{\mathcal{X}^*}$.

An important generalization of Lipschitz continuity is local-Lipschitz-continuity, a class of functions which are (Clarke) subdifferentiable.

Definition 2.8.6 [Local-Lipschitz-Continuity]**.**

A function $f : \mathcal{X} \to \mathcal{Y}$ is said to be ℓ_f-**locally Lipschitz-continuous** iff there exists $\varepsilon > 0$ s.t. for all $x \in \mathcal{X}$, f is ℓ_f-Lipschitz continuous on $\mathcal{B}_\varepsilon(x)$.

We note that any locally Lipschitz continuous function is continuous, but not vice versa.

An important refinement of continuous differentiability which has been used extensively in prior work (see, for instance, Daskalakis et al. (2020b)) is the notion of Lipschitz-smoothness, which requires the gradient of a function to Lipschitz-smooth.

Definition 2.8.7 [Lipschitz-Smoothness]**.**

A function $f : \mathcal{X} \to \mathcal{Y}$ is said to be λ-**Lipschitz-smooth on** $\mathcal{S} \subseteq \mathcal{X}$ iff for all $x_1, x_2 \in \mathcal{X}$

$$\|\nabla f(x_1) - \nabla f(x_2)\|_{\mathcal{Y}} \leq \lambda \|x_1 - x_2\|_{\mathcal{X}} , \tag{2.13}$$

When $\mathcal{S} = \mathcal{X}$, then f is simply said to be λ-**Lipschitz-smooth**.

Convexity Properties of Functions

Definition 2.8.8 [Convex/Concave Function]**.**

Consider a function $f : \mathcal{X} \to \mathbb{R}$.

f is **convex** iff for all $\lambda \in [0, 1]$ and $x, x' \in \mathcal{X}$,

$$f(\lambda x + (1 - \lambda)x') \leq \lambda f(x) + (1 - \lambda) f(x') \ .$$

f is **concave** iff for all $\lambda \in [0,1]$ and $x, x' \in \mathcal{X}$,

$$f(\lambda x + (1-\lambda)x') \geq \lambda f(x) + (1-\lambda) f(x') \ .$$

f is said to be **affine** iff it is both convex and concave.

Note that, a function f is convex iff $-f$ is concave.

When f is convex, its subdifferential correspondence $\mathcal{D}f$ is called the **convex subdifferential** and is given by $\mathcal{D}f(x) \doteq \{x^* \in \mathcal{X}^* \mid f(x') \geq f(x) + \langle x^*, x' - x \rangle, \forall x' \in \mathcal{X}\}$ (see Theorem 25.6 of Ralph Tyrell (1997)). Similarly, when f is concave, its Clarke subdifferential correspondence $\mathcal{D}f$ is called the **concave subdifferential** and is given by $\mathcal{D}f(x) \doteq \{x^* \in \mathcal{X}^* \mid f(x') \leq f(x) + \langle x^*, x' - x \rangle, \forall x' \in \mathcal{X}\}$.

Definition 2.8.9 [Quasiconvex/Quasiconcave Functions]**.**
Consider a function $f : \mathcal{X} \to \mathbb{R}$.

f is **quasiconvex** iff for all $\lambda \in (0,1)$ and $x, x' \in \mathcal{X}$,

$$f(\lambda x + (1-\lambda)x') \leq \max\{f(x), f(x')\} \ . \tag{2.14}$$

f is **quasiconcave** iff for all $\lambda \in (0,1)$ and $x, x' \in \mathcal{X}$,

$$f(\lambda x + (1-\lambda)x') \geq \min\{f(x), f(x')\} \ . \tag{2.15}$$

In addition, a function f is quasiconvex iff $-f$ is quasiconcave.

We note that a function f is quasiconvex iff its **sublevel sets**, i.e., the set $\{x \in \mathcal{X} \mid f(x) \leq \alpha\}$ for all $\alpha \in \mathbb{R}$, is convex. Similarly, a function f is quasiconcave iff its **superlevel set**, i.e., the set $\{x \in \mathcal{X} \mid f(x) \geq \alpha\}$ for all $\alpha \in \mathbb{R}$, is convex. Quasiconvex and quasiconcave functions are very useful in representing a class of continuous and convex sets or correspondences. In particular, given metric spaces $(\mathcal{X}, d_{\mathcal{X}})$, $(\mathcal{Y}, d_{\mathcal{Y}})$, and some continuous quasiconvex functions $g_1, \ldots, g_l : \mathcal{X} \times \mathcal{Y}$, then the correspondence $\mathcal{R}(y) \doteq \{x \in \mathcal{X} \mid g_i(x, y) \leq 0, i \in [l]\}$ is continuous and convex (see Theorem 5.9 of Rockafellar and Wets (2009)).

Going beyond classes of convex/concave functions, this thesis will make use of notions of weakly-convex/concave functions. The class of weakly-convex/concave functions were first introduced to the optimization literature in English by Nurminskii (1973), and have become a class of functions of great interest in the optimization literature is the recent years (see, for instance, Davis et al. (2018); Davis and Drusvyatskiy (2019); Lin et al. (2020)).

Definition 2.8.10 [Weakly-Convex/Weakly-Concave Functions]**.**

Consider a function $f : \mathcal{X} \to \mathbb{R}$.

f is μ-**weakly-convex** with **modulus of convexity** $\mu \in \mathbb{R}$, iff $x \mapsto f(x) + \mu/2\|x\|^2$ is convex. If $\mu < 0$, then f is said to be $(-\mu)$-**strongly-convex**.[7]

f is ν-**weakly-concave** with **modulus of concavity** $\nu \in \mathbb{R}$, iff $x \mapsto f(x) - \mu/2\|x\|^2$ is concave. If $\mu < 0$, then f is said to be $(-\mu)$-**strongly-convex**.

Remark 2.8.1 [Examples of Weakly-Convex/Concave Functions]**.**

Naturally, the class of weakly-convex (respectively, weakly-concave) functions generalizes strongly-convex (respectively, strongly-concave) functions, and convex (resp. concave) functions.

More importantly, however, any ℓ-smooth function is both ℓ-weakly-convex and ℓ-weakly-concave (Davis and Drusvyatskiy, 2019). Thus, in some sense, Lipschitz-smooth functions can be interpreted as being weakly-affine (i.e., both weakly-convex and weakly-concave). Nonetheless, despite Lipschitz-smooth functions being a very restricted subset of the class of weakly-convex (respectively, weakly-concave) functions, they contain a very large class of non-convex and differentiable functions. In fact, the class of Lipschitz-smooth functions on a non-empty and compact domain contains, among others, all twice-continuously differentiable functions. We refer the reader to Section 2.1 of Davis and Drusvyatskiy (2019) and Section 4 of Vial (1983) for additional results and discussions on weak-convexity and weak-concavity.

Remark 2.8.2 [Subdifferential of Weakly-Convex/Concave Functions]**.**

As convex (respectively, concave) functions are locally Lipschitz continuous, so is their difference, which in turn implies subdifferentiability by Theorem 3.1. of Clarke (2007).

[7] See, for instance, Section 9.1.2 of Boyd et al. (2004) for further characterizations.

Thus, any weakly-convex (respectively, weakly-concave) function f is subdifferentiable as it can be rewritten as the difference of two convex functions $f + {}^{\mu}\!/_{2}\|\cdot\|^2$ and ${}^{\mu}\!/_{2}\|\cdot\|^2$ (respectively, concave functions $f - {}^{\mu}\!/_{2}\|\cdot\|^2$ and ${}^{\mu}\!/_{2}\|\cdot\|^2$).

In addition, for any μ-weakly-convex (respectively, ν-weakly-concave) function f, its subdifferential correspondence is given as

$$\mathcal{D}f(x) \doteq \mathcal{D}[f(x) + \frac{\mu}{2}\|x\|^2] - \mu x \quad \left(\text{respectively, } \mathcal{D}f(x) = \mathcal{D}[f(x) - \frac{\nu}{2}\|x\|^2] + \mu x\right) \quad ,$$

where $\mathcal{D}[f(x) + {}^{\mu}\!/_{2}\|x\|^2]$ (respectively, $\mathcal{D}[f(x) - {}^{\nu}\!/_{2}\|x\|^2]$ is the convex (respectively, concave) subdifferential, since $f(\cdot) + {}^{\mu}\!/_{2}\|\cdot\|^2$ is convex (respectively, $\mathcal{D}[f(x) - {}^{\nu}\!/_{2}\|x\|^2]$ is concave).

An implication of this remark is that, as a subgradient of a convex/concave function can be computed (or in the worst case approximated) easily (see, for instance Bertsekas (2011)) via convex/concave subdifferential calculus rules, a subgradient of a weakly-convex/weakly-concave function can also be computed easily. *As such, as is standard in the literature (see for instance Lin et al. (2020)), we will take the number of subgradient evaluations as the primitive operation of our computational complexity results.*

In addition, for weakly-convex and weakly-concave functions, we have the following characterization introduced by Davis and Drusvyatskiy (see Lemma 2.1 of Davis and Drusvyatskiy (2019)) whose proof we provide for completness. Note that as a function f is μ-weakly-convex iff $-f$ is μ-weakly-concave, an analogous characterization holds for weakly-concave functions as well.

Lemma 2.8.1 [Characterization of Weakly-Convex-Functions]**.**
Consider a function $f : (\mathcal{X}, \|\cdot\|) \to \mathbb{R}$, and a modulus of convexity $\mu \in \mathbb{R}$. The following statements are equivalent:

1. f is μ-weakly convex
2. The μ**-weak secant inequality** holds:

$$f(\lambda x + (1-\lambda)x') \leq \lambda f(x) + (1-\lambda)f(x') + \frac{\mu\lambda(1-\lambda)}{2}\|x - x'\|^2 \qquad (2.16)$$

3. The μ-**weak subgradient inequality** holds, i.e.,

$$f(x') \geq f(x) + \left\langle \partial f(x), x' - x \right\rangle - \mu/2 \|x - x'\|^2, \qquad \forall x, x' \in \mathcal{X}, \partial f(x) \in \mathcal{D} f(x)$$

4. The subdifferential map is μ-**weakly monotone**, i.e.,

$$\left\langle \partial f(x') - \partial f(x), x' - x \right\rangle \geq -\mu \|x - x'\|^2 \qquad \forall \partial f(x) \in \mathcal{D} f(x), \partial f(x') \in \mathcal{D} f(x')$$

Proof of Lemma 2.8.1

Fix $\mu \in \mathbb{R}$, and let $f : \mathcal{X} \to \mathbb{R}$ be a μ-weakly-convex function. We then have:

$(1) \equiv (2)$:
From the definition of μ-weak convexity, we have for all $x, x' \in \mathcal{X}$ and $\lambda \in [0, 1]$:

$$\begin{aligned} f(\lambda x + (1-\lambda)x') + \mu/2 \|\lambda x + (1-\lambda)x'\|^2 &\leq \lambda[f(x) + \mu/2 \|x\|^2] + (1-\lambda)[f(x') + \mu/2 \|x'\|^2] \\ &\leq \lambda f(x) + (1-\lambda) f(x') + \lambda\mu/2 \|x\|^2 + (1-\lambda)\mu/2 \|x'\|^2 \ . \end{aligned}$$

Re-organizing the expression, we get:

$$f(\lambda x + (1-\lambda)x') \leq \lambda f(x) + (1-\lambda) f(x') + \lambda\mu/2 \|x\|^2 + (1-\lambda)\mu/2 \|x'\|^2 - \mu/2 \|\lambda x + (1-\lambda)x'\|^2 \quad (2.17)$$

Now, notice that we have for all $x, x' \in \mathcal{X}$ and $\lambda \in [0, 1]$:

$$\begin{aligned} &\lambda\mu/2 \|x\|^2 + (1-\lambda)\mu/2 \|x'\|^2 - \mu/2 \|\lambda x + (1-\lambda)x'\|^2 \\ &= \lambda\mu/2 \|x\|^2 + (1-\lambda)\mu/2 \|x'\|^2 - \mu\lambda^2/2 \|x\|^2 - \lambda(1-\lambda)\mu \left\langle x, x' \right\rangle - (1-\lambda)^2\mu/2 \|x'\|^2 \\ &= \lambda\mu/2 \|x\|^2 - \mu\lambda^2/2 \|x\|^2 + (1-\lambda)\mu/2 \|x'\|^2 - (1-\lambda)^2\mu/2 \|x'\|^2 - \lambda(1-\lambda)\mu \left\langle x, x' \right\rangle \\ &= \lambda(1-\lambda)\mu/2 \|x\|^2 + \lambda(1-\lambda)\mu/2 \|x'\|^2 - \lambda(1-\lambda)\mu \left\langle x, x' \right\rangle \\ &= \frac{\lambda(1-\lambda)\mu}{2} \|x - x'\|^2 \end{aligned}$$

Hence, plugging the above into Equation (2.17), we get for all $x, x' \in \mathcal{X}$ and $\lambda \in [0, 1]$:

$$f(\lambda x + (1-\lambda)x') \leq \lambda f(x) + (1-\lambda) f(x') + \frac{\lambda(1-\lambda)\mu}{2} \|x - x'\|^2$$

As all the inequalities are tight, the implication holds both ways.

$(2) \equiv (3)$

Re-organizing the terms in Equation (2.16), we have for all $x, x' \in \mathcal{X}$ and $\lambda \in [0, 1]$:

$$f(\lambda x + (1-\lambda)x') \leq \lambda f(x) + (1-\lambda) f(x') + \lambda(1-\lambda)\mu/2 \|x - x'\|^2$$

$$f(x' + \lambda(x - x')) \leq \lambda f(x) + (1-\lambda) f(x') + \lambda(1-\lambda)\mu/2 \|x - x'\|^2$$

$$f(x' + \lambda(x - x')) - f(x') \leq \lambda \left[f(x) - f(x') \right] + \lambda(1-\lambda)\mu/2 \|x - x'\|^2$$

Dividing both sides of the inequality by λ, for all $x, x' \in \mathcal{X}$ and $\lambda \in [0,1]$ we have:

$$\frac{f(x' + \lambda(x - x')) - f(x)}{\lambda} \leq f(x) - f(x') + {}^{(1-\lambda)\mu}/_{2}\|x - x'\|^2$$

Now, taking the limit as $\lambda \to 0$, we have for all $x, x' \in \mathcal{X}$ the μ-subgradient inequality for all possible subgradients:

$$\lim_{\lambda \to 0} \frac{f(x' + \lambda(x' - x)) - f(x')}{\lambda} \leq f(x) - f(x') + {}^{(1-\lambda)\mu}/_{2}\|x - x'\|^2$$

Note that left hand side is the definition of the Gâteau derivative. Hence, by varying $x \in \mathcal{X}$, we obtain all possible limits points of the Gâteau derivative. As this set of limit points is closed and convex, it is equal to its convex hull, which is the definition of the Clarke subdifferential. In addition, once again, as we have applied no inequalities, the bounds are tight and the implication holds both ways.

$(3) \equiv (4)$:

By the μ-subgradient inequality, we have the two following relations:

$$f(x') \geq f(x) + \left\langle \partial f(x), x' - x \right\rangle - {}^{\nu}/_{2}\|x - x'\|^2, \qquad \forall x, x' \in \mathcal{X}, \partial f(x) \in \mathcal{D}f(x)$$

$$f(x) \geq f(x') + \left\langle \partial f(x'), x - x' \right\rangle - {}^{\nu}/_{2}\|x - x'\|^2, \qquad \forall x, x' \in \mathcal{X}, \partial f(x') \in \mathcal{D}f(x')$$

Subtracting the first inequality from the second, and re-organizing terms, we obtain the μ-weak-monotonicity condition. The reverse direction follows in the same way.

2.9 Constrained Optimization Background

2.9.1 The Primal Problem

Consider any Euclidean metric space $(\mathbb{R}^n, d)$. A **(constrained) optimization problem** $\mathcal{C} \doteq (n, l, \mathcal{X}, f, \boldsymbol{g})$, denoted $(\mathcal{X}, f, \boldsymbol{g})$ when n and l are clear from context, consists of an **objective function** $f : \mathbb{R}^n \to \mathbb{R}$, $l \in \mathbb{N}$ **constraint functions** $g_1, \ldots, g_l : \mathbb{R}^n \to \mathbb{R}$, and a basic feasible set $\mathcal{X} \subseteq \mathbb{R}^n$, which together define the following maximization problem called the **Primal Problem**:[8]

Primal Problem

$$\max_{\boldsymbol{x} \in \mathcal{X}} f(\boldsymbol{x}) \tag{2.18}$$

$$\text{constrained by} \qquad g_i(\boldsymbol{x}) \geq 0 \qquad \forall i \in [l] \tag{2.19}$$

For convenience, we define $\boldsymbol{g} \doteq (g_i)_{i=1}^l : \mathbb{R}^n \to \mathbb{R}^l$, and consistent with the partial order we defined on Euclidean vector spaces, we will often write $\boldsymbol{g}(\boldsymbol{x}) \geq \boldsymbol{0}$ to mean for all $i \in [l]$, $g_i(\boldsymbol{x}) \geq 0$.

We define the **feasible set** $\mathrm{feas}(\mathcal{X}, f, \boldsymbol{g})$ of the optimization problem $(\mathcal{X}, f, \boldsymbol{g})$ as:

$$\mathrm{feas}(\mathcal{X}, f, \boldsymbol{g}) \doteq \{\boldsymbol{x} \in \mathcal{X} \mid \boldsymbol{g}(\boldsymbol{x}) \geq \boldsymbol{0}\}$$

A point $\boldsymbol{x} \in \mathbb{R}^n$ is said to be **feasible** if it is an element of the feasible set $\mathrm{feas}(\mathcal{X}, f, \boldsymbol{g})$. Note that the primal problem can be restated as $\max_{\boldsymbol{x} \in \mathrm{feas}(\mathcal{X}, f, \boldsymbol{g})} f(\boldsymbol{x})$. As such, by Weierstrass' Extreme Value Theorem, under the following assumption, which we assume throughout this section, a solution to the primal problem exists:

Assumption 2.9.1 [Existence of Solution]**.**

Consider an optimization problem $(\mathcal{X}, f, \boldsymbol{g})$. Assume

1. $f : \mathcal{X} \to \mathbb{R}$ is continuous
2. $\mathrm{feas}(\mathcal{X}, f, \boldsymbol{g})$ is non-empty and compact

[8] For convenience, we focus on maximization problems. This convention is without loss of generality, since any maximization problem can be recast as a minimization by negating the objective functions.

We note that Part 2 of Assumption 2.9.1, can be guaranteed under the assumption that $\boldsymbol{g}$ is continuous, and the feasible set $\text{feas}(\mathcal{X}, f, \boldsymbol{g})$ is non-empty.

2.9.2 The Lagrangian and the Dual Problem

For any optimization problem $(\mathcal{X}, f, \boldsymbol{g})$, we define the **Lagrangian (function)** $\ell : \mathbb{R}^n \times \mathbb{R}^l_+ \to \mathbb{R}$ as:

$$\ell(\boldsymbol{x}, \boldsymbol{\lambda}) \doteq f(\boldsymbol{x}) + \sum_{i=1}^{l} \lambda_i g_i(\boldsymbol{x})$$

where $\boldsymbol{\lambda} \doteq (\lambda_1, \ldots, \lambda_l) \in \mathbb{R}^l_+$ are called **slack variables** (or **KKT multipliers**). These variables are called slack variables as they relax the constrained problem to an unconstrained one, and then, by selecting their values wisely, we obtain a function whose maximum over $\mathcal{X}$ corresponds exactly to that of the primal problem. More formally, taking the infimum of the Lagrangian over the respective domains of the slack variables, we have:

$$\begin{aligned}
\inf_{\boldsymbol{\lambda} \geq \mathbf{0}} \ell(\boldsymbol{x}, \boldsymbol{\lambda}) &= \inf_{\boldsymbol{\lambda} \geq \mathbf{0}} \left[f(\boldsymbol{x}) + \sum_{i=1}^{l} \lambda_i g_i(\boldsymbol{x}) \right] \\
&= f(\boldsymbol{x}) + \inf_{\boldsymbol{\lambda} \geq \mathbf{0}} \sum_{i=1}^{l} \lambda_i g_i(\boldsymbol{x}) \\
&= f(\boldsymbol{x}) + \sum_{i=1}^{l} \inf_{\lambda_i \geq 0} \lambda_i g_i(\boldsymbol{x})
\end{aligned} \tag{2.20}$$

Note that for all $i \in [l]$,

$$\inf_{\lambda_i \geq 0} \lambda_i g_i(\boldsymbol{x}) = \begin{cases} 0 & \text{if } g_i(\boldsymbol{x}) \geq 0 \\ \infty & \text{Otherwise} \end{cases}$$

Plugging back into Equation (2.20), we then have:

$$\inf_{\boldsymbol{\lambda} \geq \mathbf{0}} \ell(\boldsymbol{x}, \boldsymbol{\lambda}) = \begin{cases} f(\boldsymbol{x}) & \text{if } \boldsymbol{g}(\boldsymbol{x}) \leq \mathbf{0} \\ \infty & \text{Otherwise} \end{cases}$$

That is, by taking the infimum of the Lagrangian over the slack variables $\boldsymbol{\lambda}$, we obtain a function where for all feasible points $\boldsymbol{x} \in \text{feas}(\mathcal{X}, f, \boldsymbol{g})$, the value of function coincides with the value of the objective f, and for all infeasible points $\boldsymbol{x}' \notin \text{feas}(\mathcal{X}, f, \boldsymbol{g})$, the value of the function is ∞. As a result, we have:

$$\max_{\boldsymbol{x} \in \text{feas}(\mathcal{X}, f, \boldsymbol{g})} f(\boldsymbol{x}) = \max_{\boldsymbol{x} \in \mathcal{X}} \inf_{\boldsymbol{\lambda} \geq \boldsymbol{0}} \ell(\boldsymbol{x}, \boldsymbol{\lambda})$$

The above equality suggests that if we could switch the order of the max and inf on the right hand-side, we could re-express primal problem, as a minimization problem called the **dual problem**. To this end, the **Lagrangian dual function** $\ell^* : \mathbb{R}^l_+ \to \bar{\mathbb{R}}$ is defined as:

$$\ell^*(\boldsymbol{\lambda}) = \sup_{\boldsymbol{x} \in \mathcal{X}} \ell(\boldsymbol{x}, \boldsymbol{\lambda}) \tag{2.21}$$

The **dual problem** associated with any optimization problem $(\mathcal{X}, f, \boldsymbol{g})$ is then defined as:

$$\inf_{\boldsymbol{\lambda}} \ell^*(\boldsymbol{\lambda}) \tag{2.22}$$

$$\text{constrained by} \qquad \boldsymbol{\lambda} \geq \boldsymbol{0} \tag{2.23}$$

A tuple of slack variables $\boldsymbol{\lambda} \in \mathbb{R}^l$ are said to be **feasible** iff $\boldsymbol{\lambda} \geq \boldsymbol{0}$.

We say that **weak duality holds** iff $\max_{\boldsymbol{x} \in \mathcal{X}} \inf_{\boldsymbol{\lambda} \geq \boldsymbol{0}} \ell(\boldsymbol{x}, \boldsymbol{\lambda}) \leq \inf_{\boldsymbol{\lambda} \geq \boldsymbol{0}} \sup_{\boldsymbol{x} \in \mathcal{X}} \ell(\boldsymbol{x}, \boldsymbol{\lambda})$. We say that **strong duality holds** iff $\max_{\boldsymbol{x} \in \mathcal{X}} \min_{\boldsymbol{\lambda} \geq \boldsymbol{0}} \ell(\boldsymbol{x}, \boldsymbol{\lambda}) = \min_{\boldsymbol{\lambda} \geq \boldsymbol{0}} \max_{\boldsymbol{x} \in \mathcal{X}} \ell(\boldsymbol{x}, \boldsymbol{\lambda})$. Note that for strong duality the infimum and supremum are replaced by a minimum and maximum since under Assumption 2.9.1 $\sup_{\boldsymbol{x} \in \mathcal{X}} \inf_{\boldsymbol{\lambda} \geq \boldsymbol{0}} \ell(\boldsymbol{x}, \boldsymbol{\lambda})$ is well-defined, and as such for equality to hold the supremum and infimum must be well-defined. Without the need for any additional assumptions, we can show that weak duality holds for any optimization problem.

Theorem 2.9.1 [Weak Duality]**.**
Given an optimization problem $(\mathcal{X}, f, \boldsymbol{g})$, under Assumption 2.9.1, weak duality holds.

Proof

$$\begin{aligned}
\inf_{\boldsymbol{\lambda} \geq \boldsymbol{0}} \ell(\boldsymbol{x}, \boldsymbol{\lambda}) &\leq \ell(\boldsymbol{x}, \boldsymbol{\lambda}') && \forall \boldsymbol{x} \in \mathcal{X}, \boldsymbol{\lambda}' \geq \boldsymbol{0} \\
\sup_{\boldsymbol{x} \in \mathcal{X}} \inf_{\boldsymbol{\lambda} \geq \boldsymbol{0}} \ell(\boldsymbol{x}, \boldsymbol{\lambda}) &\leq \sup_{\boldsymbol{x} \in \mathcal{X}} \ell(\boldsymbol{x}, \boldsymbol{\lambda}') && \forall \boldsymbol{\lambda}' \geq \boldsymbol{0} \\
\sup_{\boldsymbol{x} \in \mathcal{X}} \inf_{\boldsymbol{\lambda} \geq \boldsymbol{0}} \ell(\boldsymbol{x}, \boldsymbol{\lambda}) &\leq \inf_{\boldsymbol{\lambda} \geq \boldsymbol{0}} \sup_{\boldsymbol{x} \in \mathcal{X}} \ell(\boldsymbol{x}, \boldsymbol{\lambda})
\end{aligned}$$

In contrast, to ensure that strong duality holds, we have to restrict the class of optimization problems considerably, namely to **regular** convex optimization problems (i.e., optimization problems where f and $\boldsymbol{g}$ are convex, and for which a constraint qualification such as Slater's condition is satisfied). For example, we can make the following additional assumption:

Assumption 2.9.2 [Convex optimization and Slater's condition]**.**
Consider an optimization problem $(\mathcal{X}, f, \boldsymbol{g})$. Suppose that:

1. f is concave
2. $\boldsymbol{g}$ is concave
3. (Slater's condition) There exists an feasible relative interior point $\widehat{\boldsymbol{x}} \in \operatorname{relint}(\operatorname{feas}(\mathcal{X}, f, \boldsymbol{g}))$, i.e., $\widehat{\boldsymbol{x}} \in \operatorname{int}(\mathcal{X})$ and for all $c \in [l]$, if g_c is affine then $g_c(\widehat{\boldsymbol{x}}) \geq 0$, and $g_c(\widehat{\boldsymbol{x}}) > 0$ otherwise.

As the proof of the strong duality theorem is more involved, we state the theorem without proof and refer the reader to page 234 of Boyd et al. (2004). Briefly, the theorem is usually proven via the Separating Hyperplane Theorem.

Theorem 2.9.2 [Strong duality via Slater's condition]**.**
Consider an optimization problem $(\mathcal{X}, f, \boldsymbol{g})$ and suppose that Assumptions 2.9.1 and 2.9.2 are satisfied, then strong duality holds.

When strong duality holds, for any solution $\boldsymbol{x}^* \in \mathcal{X}$ of the primal problem, we are guaranteed the existence of some associated optimal slack variables $\boldsymbol{\lambda}^* \in \mathbb{R}^l_+$ s.t. $\boldsymbol{\lambda}^* \in \arg\min_{\boldsymbol{\lambda} \in \mathbb{R}^l_+} \ell(\boldsymbol{x}^*, \boldsymbol{\lambda})$, which are solutions to the dual problem. More importantly, under strong duality, we can derive necessary conditions that characterize the any tuple $(\boldsymbol{x}^*, \boldsymbol{\lambda}^*) \in \mathcal{X} \times \mathbb{R}^l_+$ of primal and dual problem solutions, known as the Kahn-Karush-Tucker (KKT) conditions (Kuhn and Tucker, 1951). For convenience, we state the next theorem under the simplifying assumption that $\mathcal{X} \doteq \mathbb{R}^n$.[9]

[9]When $\mathcal{X}$ can be represented by finitely many (in)equality constraints, this assumption is without loss of generality, as $\mathcal{X}$ can be represented by the (in)equality constraint functions $\boldsymbol{g}$. Alternatively, if a relative interior solution to the primal problem exists, then the assumption is also without loss of generality. In addition, the KKT Theorem can be generalized to arbitrary $\mathcal{X}$; however, as this more general characterization will not be used in this thesis, we omit it here.

Theorem 2.9.3 [Karush–Kuhn–Tucker theorem]**.**

Consider a Euclidean normed vector space $(\mathbb{R}^n, \|\cdot\|)$, and an associated optimization problem $(\mathcal{X}, f, \boldsymbol{g})$. Suppose that Assumption 2.9.1 and Assumption 2.9.2 hold, and that in addition, f and $\boldsymbol{g}$ are locally Lipschitz continuous. Then, any tuple $(\boldsymbol{x}^*, \boldsymbol{\lambda}^*) \in \mathcal{X} \times \mathbb{R}^l$ of primal and dual problem solutions must satisfy the following conditions:

1. (Stationarity) $\mathbf{0} \in \mathcal{D}f(\boldsymbol{x}^*) + \sum_{i=1}^{l} \lambda_i^* \mathcal{D}g_i(\boldsymbol{x}^*)$
2. (Complementary Slackness) for all $i \in [l]$, $\lambda_i^* g_i(\boldsymbol{x}^*) = 0$
3. (Primal Feasibility) $\boldsymbol{g}(\boldsymbol{x}^*) \geq \mathbf{0}$
4. (Dual Feasibility) $\boldsymbol{\lambda} \geq \mathbf{0}$

In the above theorem, if we assume that Assumption 2.9.2, holds, then the conditions are also sufficient.

Often, we also are interested in understanding the convexity properties of set of solutions to the primal and dual problems. A proof of the following result can be obtained by combining the results in Chapter 6, Section 3 of Berge (1997) and in Proposition 4.1 of Kyparisis (1985).

Theorem 2.9.4 [Properties of the primal solution set]**.**

Given an optimization problem $(\mathcal{X}, f, \boldsymbol{g})$. Assume

1. f is continuous, quasiconcave
2. $\mathrm{feas}(\mathcal{X}, f, \boldsymbol{g})$ is non-empty, compact, and convex.

Then, the set of solutions to the primal problem is non-empty, compact, and convex.

If we instead assume Assumption 2.9.2, we obtain a characterization of both the primal and the dual solution sets (see Theorem 5 of Rockafellar (1971)).

Theorem 2.9.5 [Properties of the saddle point solution set]**.**

Consider an optimization problem $(\mathcal{X}, f, \boldsymbol{g})$. Suppose that Assumption 2.9.1 and Assumption 2.9.2 hold. Then, the set of solutions to the primal and dual problem is non-empty, compact, and convex.

2.9.3 Parametric Constrained Optimization

Many problems of interest can be posed as parametric constrained optimization problems, i.e., optimization problems $(\mathcal{X}, f, \boldsymbol{g})$ in which the objective $\boldsymbol{x} \mapsto f(\boldsymbol{x}; \boldsymbol{\theta})$ and the constraint function $\boldsymbol{x} \mapsto \boldsymbol{g}(\boldsymbol{x}; \boldsymbol{\theta})$ depend on the value of some parameter $\boldsymbol{\theta} \in \Theta$.

Consider metric spaces $(\mathbb{R}^n, d_{\mathcal{X}})$ and $(\mathbb{R}^d, d_\Theta)$. A **parametric (constrained) optimization problem** $(n, d, l, \Theta, \mathcal{X}, f, \boldsymbol{g})$, denoted $(\Theta, \mathcal{X}, f, \boldsymbol{g})$ when n, d, and l are clear from context, consists of a **basic feasible set** $\mathcal{X} \subseteq \mathbb{R}^n$, **set of parameters** $\Theta \subseteq \mathbb{R}^d$, a **parametric objective function** $f : \mathbb{R}^n \times \mathbb{R}^d \to \mathbb{R}$, and an **inequality constraint function** $\boldsymbol{g} : \mathbb{R}^n \times \mathbb{R}^d \to \mathbb{R}^l$, which for all $\boldsymbol{\theta} \in \Theta$ define the following maximization problem:

$$\max_{\boldsymbol{x} \in \mathcal{X}} f(\boldsymbol{x}; \boldsymbol{\theta}) \tag{2.24}$$

$$\text{constrained by} \qquad \boldsymbol{g}(\boldsymbol{x}; \boldsymbol{\theta}) \geq \boldsymbol{0} \tag{2.25}$$

When faced with a parametric optimization problem, we are often interested in understanding properties of the **marginal function** $f^*(\boldsymbol{\theta}) \doteq \max_{\boldsymbol{x} \in \mathcal{X} : \boldsymbol{g}(\boldsymbol{x}) \geq \boldsymbol{0}} f(\boldsymbol{x}; \boldsymbol{\theta})$ and the **solution correspondence** $\mathcal{X}^*(\boldsymbol{\theta}) \doteq \arg\max_{\boldsymbol{x} \in \mathcal{X} : \boldsymbol{g}(\boldsymbol{x}) \geq \boldsymbol{0}} f(\boldsymbol{x}; \boldsymbol{\theta})$. The maximum theorem, also known as Berge's maximum theorem (Berge, 1997), characterizes properties of the marginal function and solution correspondence. For convenience, define the **constraint correspondence** $\mathcal{C}(\boldsymbol{\theta}) \doteq \{\boldsymbol{x} \in \mathcal{X} \mid \boldsymbol{g}(\boldsymbol{x}; \boldsymbol{\theta}) \geq \boldsymbol{0}\}$.

Theorem 2.9.6 [Maximum Theorem for Continuity]**.**
Consider a **parametric optimization problem** $(\Theta, \mathcal{X}, f, \boldsymbol{g})$. If the constraint correspondence $\mathcal{C}$ is continuous and non-empty, compact-valued, and f is continuous, then the following hold:

1. f^* is continuous, and
2. $\mathcal{X}^*$ is upper hemicontinuous, and non-empty, and compact-valued.

We note that the continuity cf the constraint correspondence $\mathcal{C}$ can be guaranteed under either of the following assumptions by Theorem 5.9 and Example 5.10 of Rockafellar and Wets (2009):

Assumption 2.9.3 [Continuity via quasiconvex representation]**.**
Consider a correspondence $\mathcal{C}(\boldsymbol{\theta}) \doteq \{\boldsymbol{x} \in \mathcal{X} \mid \boldsymbol{g}(\boldsymbol{x};\boldsymbol{\theta}) \geq \boldsymbol{0}\}$, and suppose that:

1. $\mathcal{X}$ is compact
2. $\boldsymbol{g}$ is continuous and quasiconcave

Assumption 2.9.4 [Continuity via Slater's condition]**.**
Consider a correspondence $\mathcal{C}(\boldsymbol{\theta}) \doteq \{\boldsymbol{x} \in \mathcal{X} \mid \boldsymbol{g}(\boldsymbol{x};\boldsymbol{\theta}) \geq \boldsymbol{0}\}$, and suppose that

1. $\mathcal{X}$ is compact
2. $\boldsymbol{g}$ is continuous and concave
3. (Slater's condition) For all $\boldsymbol{\theta} \in \Theta$, there exists a feasible relative interior point $\widehat{\boldsymbol{x}} \in \mathrm{relint}(\mathcal{C}(\boldsymbol{\theta}))$, i.e., $\widehat{\boldsymbol{x}} \in \mathrm{int}(\mathcal{X})$ and for all $c \in [l]$, if $\boldsymbol{x} \mapsto g_c(\boldsymbol{x},\boldsymbol{\theta})$ is affine then $g_c(\widehat{\boldsymbol{x}},\boldsymbol{\theta}) \geq 0$, and $g_c(\widehat{\boldsymbol{x}},\boldsymbol{\theta}) > 0$ otherwise.

Going further than continuity, we might more generally be interested in understanding the convexity and concavity properties of the marginal function and solution correspondence. As a corollary, of Theorem 2.9.4 we have the following convexity characterization of the image of the solution correspondence.

Corollary 2.9.1 [Convex-valued solution correspondence]**.**
Consider a **parametric optimization problem** $(\Theta, \mathcal{X}, f, \boldsymbol{g})$. If the constraint correspondence $\mathcal{C}$ is non-empty-, compact-, and convex-valued, and for all $\boldsymbol{\theta} \in \Theta$, $\boldsymbol{x} \mapsto f(\boldsymbol{x};\boldsymbol{\theta})$ is continuous and quasiconcave, then $\mathcal{X}^*$ is non-empty-, and compact-, and convex-valued.

The following theorem provides sufficient conditions for the concavity of the marginal function and convexity of the solution correspondence (see Theorem 2.1 and 3.1 of Kyparisis and Fiacco (1987)).

Theorem 2.9.7 [Convexity/concavity of the marginal function]**.**

Consider a **parametric optimization problem** $(\Theta, \mathcal{X}, f, \boldsymbol{g})$. If the constraint correspondence $\mathcal{C}$ is convex, and f is concave, then f^* and $\mathcal{X}^*$ are concave.

Part I

Variational Inequalities and Walrasian Economies

Chapter 3

Scope and Motivation

3.1 Scope

Part I of this thesis is divided into three chapters. In Chapter 4, after reviewing background material on variational inequalities, we[1] will introduce two new type of method with polynomial-time convergence guarantees. The first type of methods will be a family of first-order methods which we call the mirror extragradient method. We will prove that this method converges to a strong solution of any variational inequality for which a weak solution exists. Further, in the absence of a weak solution, we will prove local convergence to a strong solution when the algorithm's first iterate is initialized close enough to a local weak solution. As first order methods are not guaranteed to converge beyond settings where a (local) weak solution exists, we will then turn our attention to a class of second-order methods known as merit function methods. In particular, we will introduce the primal mirror descent method, which we will show is guaranteed to converge to a local minimum of the regularized primal gap function of any Lipschitz-smooth variational inequality.

In Chapter 5, after reviewing background material on Walrasian economies, we will show that the set of Walrasian equilibria of any Walrasian economy is equal to set of strong solutions of an associated variational inequality. In addition, we will show that the mirror

[1]The works in Part I of this thesis were a collaboration with Amy Greenwald, with Sadie Zhao additionally contributing to the verification of Chapter 6.

gradient method applied to this variational inequality is equivalent to solving a Walrasian economy via a well-known price-adjustement process known as *tâtonnement*. Furthermore, by running the mirror extragradient method on this variational inequality, we obtain a new family of price adjustment processes called the mirror *extratâtonnement* process, for which we show polynomial-time convergence to a Walrasian equilibrium in a large class of Walrasian economies. Finally, using the VI characterization of Walrasian equilibrium, we will introduce a class of merit function methods with polynomial-time convergence guarantees for an even broader class of Walrasian economies.

While the results in Chapter 5 answer a number of open questions relevant to the variational inequality framework in general, in Chapter 6, using the tools of convex optimization and consumer theory, we provide a more fine-grained convergence analysis with better guarantees for a particular *tâtonnement* process in a class of Walrasian economies used widely in practice, known as Fisher markets. In particular, we will show the sublinear convergence of entropic *tâtonnement* to a Walrasian equilibrium in homothetic Fisher markets with bounded elasticity of Hicksian demand.

3.2 Motivation

Walrasian economies (or general equilibrium models), first studied by French economist Léon Walras in 1874, are a broad mathematical framework for modeling any economic system governed by supply and demand (Walras, 1896). A Walrasian economy consists of a finite set of commodities, characterized by an excess demand function that maps values for commodities, called **prices**, to positive (respectively, negative) quantities of each commodity demanded (respectively, supplied) in excess. Walras proposed a steady-state solution of his economy, namely a Walrasian (or competitive) equilibrium, represented by a collection of per-commodity prices which is **feasible**, i.e., there is no excess demand for any commodity, and for which **Walras' law** holds, i.e., the value of the excess demand is equal to 0.

Walras did not establish conditions ensuring the existence of an equilibrium, but he did argue, albeit without conclusive evidence, that his economy would settle at a Walrasian equilibrium via a **price-adjustment process** known as ***tâtonnement***, which mimics the behavior of the **law of supply and demand**. More specifically, this process generates a sequence of prices based on prior prices and associated excess demands, updating prices at a rate equal to the excess demand (Walras, 1896; Uzawa, 1960; Arrow and Hurwicz, 1958). To motivate the relevance of *tâtonnement* to real-world economies, Walras argued that *tâtonnement* is a **natural price-adjustment process**, in the sense that if each each commodity is owned by a different seller, then each seller can update the price of its commodity without coordinating with other sellers, using information only about the excess demand of its own commodity, hence making it plausible that *tâtonnement* could explain the movement of prices in real-world economies where sellers again do not coordinate with one another.

Nearly half a century after Walras' initial foray into general equilibrium analysis, a group of academics brought together by the Cowles Commission in 1939 reinitiated a study of Walras' economic model with the purpose of bringing rigorous mathematics to the analysis of markets. One of the earliest and most important outputs of this collaborative effort was the introduction of a broad and well-justified class of Walrasian economies known as **competitive economies** (Arrow and Debreu, 1954), for which the existence of Walrasian equilibrium was established by a novel application of fixed point theorems to economics. With the question of existence thus resolved, the field subsequently turned its focus to investigating questions on the **stability** of Walrasian equilibrium i.e., which price-adjustment processes can settle at a Walrasian equilibrium and under what assumptions? (Uzawa, 1960; Balasko, 1975; Arrow and Hurwicz, 1958; Cole and Fleischer, 2008; Cheung et al., 2018; 2013; Jain et al., 2005; Codenotti et al., 2005; 2006; Chen and Teng, 2009).

Most relevant work on stability has been concerned with the convergence properties of *tâtonnement*. Beyond Walras' justification for *tâtonnement*'s relevance to real-world economies, research on *tâtonnement* in the post-world war II economics literature is motivated by the fact that it can be understood as a plausible explanation of how prices move in real-world markets (Gillen et al., 2020). Hence, if one could prove that *tâtonnement* is a **universal price-adjustment process** (i.e., a price-adjustment process that converges to a Walrasian equilibrium in *all* competitive economies), then perhaps it would be justifiable to claim real-world economies would also eventually settle at a Walrasian equilibrium.

In 1958, Arrow and Hurwicz (1958) established the convergence of a continuous-time variant of *tâtonnement* in Walrasian economies with an excess demand function satisfying the **weak axiom of revealed preferences (WARP)** (Afriat, 1967), which among others, includes Walrasian economies satisfying the **gross substitutes (GS)** condition (Arrow et al., 1959; Arrow and Hurwicz, 1960) (i.e., if the price of one commodity increases, the excess demand for the other commodities does not decrease). This result was complemented by Nikaidô and Uzawa's (Nikaidô and Uzawa, 1960) result on the convergence of a discrete-time variant of *tâtonnement* in Walrasian economies satisfying WARP—albeit without any non-asymptotic convergence guarantees. These initial results sparked hopes that *tâtonnement* could be a universal price-adjustment process.

Furthermore, as there in general there does not exist a closed-form characterization of Walrasian equilibria, these results ignited further interest in discovering algorithms to compute a Walrasian equilibrium, as *tâtonnement* could be implemented on a computer to obtain numerical approximations of Walrasian equilibria. Indeed, these early results on the stability of *tâtonnement* inspired a new line of work on **applied general equilibrium** (Scarf, 1967b;a; Scarf and Hansen, 1973; Scarf, 1982) initiated by Herbert Scarf[2], whose goal was to establish "a general method for the explicit numerical solution of the neoclassical [Walrasian economy] model" (Scarf and Hansen, 1973). The motivation behind this research agenda was a desire to predict the impact of economic policy by estimating the parameters of a parametric Walrasian economy from empirical data, and then running a comparative

[2]See, for instance, Arrow and Kehoe (1994) for a detailed exposition of Herbert Scarf's contribution to general equilibrium theory.

static analysis to compare the numerical solution of the Walrasian economy before and after the implementation of the policy.

Unfortunately, soon after initiating this research agenda, Scarf dashed all hopes that *tâtonnement* could be a universal price-adjustment process by showing that the sequence of prices generated by a continuous-time variant of *tâtonnement* can cycle ad infinitum around the Walrasian equilibrium of his eponymous competitive economy, with only three commodities and an excess demand function generated by three consumers with Leontief preferences, i.e., **the Scarf economy** (Scarf, 1960). Even more discouragingly, when applied to the Scarf economy, the prices generated by discrete-time variants of *tâtonnement* spiral away from the Walrasian equilibrium, moving further and further away from equilibrium.

Scarf's negative result seems to have discouraged further research by economists on the stability of Walrasian equilibrium (Fisher, 1975). Despite research on this question coming to a near halt, one positive outcome was achieved, on the convergence of a non-*tâtonnement* update rule known as **Smale's process** (Herings, 1997; Kamiya, 1990; van der Laan et al., 1987; Smale, 1976), which updates prices at the rate of the product of the excess demand and the inverse of its Jacobian, to a Walrasian equilibrium in competitive economies that have an excess demand with a non-singular Jacobian, including Scarf economies. Unfortunately, this convergence result for Smale's process comes with two caveats: 1) Smale's process is not a "natural" price-adjustment process, as it updates the price of each commodity using information about not only the excess demand of the commodity but also the derivative of the excess demand function with respect to each commodity in the economy, 2) convergence of discrete time-variants of Smale's process requires that the excess demand function satisfy the law of supply and demand (i.e., a monotone excess demand correspondence), which even Walrasian economies that satisfy the GS or WARP conditions need not satisfy.

Nearly half a century after these seminal analyses of competitive economies, research on the stability and efficient computation of Walrasian equilibrium is once again coming to the fore, motivated by applications of algorithms to compute Walrasian equilibrium in dynamic stochastic general equilibrium models in macroeconomics (Geanakoplos, 1990; Sargent and Ljungqvist, 2000; Taylor and Woodford, 1999; Fernández-Villaverde, 2023), and the use of algorithms such as *tâtonnement* to solve models of transactions on crypotocurrency blockchains (Leonardos et al., 2021; Liu et al., 2022; Reijsbergen et al., 2021) and load balancing over networks (Jain et al., 2013). In contrast to the prior literature on the stability of *tâtonnement*, which was primarily concerned with proving asymptotic convergence of price-adjustment processes to a Walrasian equilibrium, this line of work is also concerned with obtaining *non-asymptotic* convergence rates, and hence computing approximate Walrasian equilibria in polynomial time.

The first result in this direction is due to Codenotti et al. (2005), who introduce a discrete-time version of *tâtonnement*, and show that in exchange economies that satisfy **weak gross substitutes (WGS)** (i.e., the excess demand of any commodity *weakly* increases if the price of any other commodity increases, fixing all other prices), the *tâtonnement* process converges to an approximate Walrasian equilibrium in a number of steps which is polynomial in the inverse of the approximation factor and size of the problem. Unfortunately, soon after this positive result appeared, Papadimitriou and Yannakakis (2010) showed that it is impossible for a price-adjustment process based on excess demand to converge in polynomial time to a Walrasian equilibrium in general (i.e., Walrasian economies with a Lipschitz-continuous excess demand which satisfy homogeneity of degree 0 and Walras' law), ruling out the possibility of Smale's process (and many others), justifying the notion of Walrasian equilibrium in all competitive economies. Nonetheless, further study of the convergence of price-adjustment processes such as *tâtonnement* under stronger assumptions, or in simpler models than full-blown Arrow-Debreu competitive economies, continues, as these processes are being deployed in practice (Jain et al., 2013; Leonardos et al., 2021; Liu et al., 2022; Reijsbergen et al., 2021).[3]

[3]We refer the reader to Sections 4.2.2 to 5.3.2 for additional related work on algorithms for solving Walrasian economies and VIs.

3.3 Contributions

3.3.1 A Tractable Variational Inequality Framework for Walrasian Economies

To address the challenge brought forward by the impossibility result of Papadimitriou and Yannakakis (2010), we provide a characterization of Walrasian equilibrium using the variational inequality (VI) optimization framework. To this end, we first introduce the class of **mirror extragradient algorithms** and prove the polynomial-time convergence of this method for VIs that satisfy a computational tractability condition known as the Minty condition (Minty, 1967), and a generalization of Lipschitz-continuity known as Bregman-continuity. A Bregman-continuous (or relatively continuous (Lu, 2019)) function is one for which the change in the Euclidean distance of the function between any two points is proportional to the Bregman divergence between those two points.

With these tools in place, we then demonstrate that the set of Walrasian equilibria of **balanced economies**—those Walrasian economies with an excess demand correspondence that is homogeneous of degree 0, and satisfy weak Walras's (i.e., the value of the excess demand is less than or equal to 0 at all prices)—a class of Walrasian economies which among others includes Arrow and Debreu's competitive economies (Arrow and Debreu, 1954), is equal to the set of strong solutions of a VI that satisfies the Minty condition (Minty, 1967). With this characterization in hand, we apply the mirror extragradient algorithm to solving this VI, which gives rise to a novel natural price-adjustment process we call ***mirror extratâtonnement***.

An important property of the VI we introduce is that its search space for prices is *not* restricted to the unit simplex as it is traditionally the case for competitive economies, but rather to the unit box. This fact offers us insight into understanding how we can overcome Papadimitriou and Yannakakis's impossibility result on the exponential-time convergence of price-adjustment processes in general Walrasian economies. Papadimitriou and Yannakakis's definition of a price-adjustment process restricts prices generated by the process to lie within the unit simplex; indeed, when the search space of the VI we introduce is restricted in this way, the VI fails to satisfy the Minty condition, and is thus computationally intractable. This suggests that relaxing the requirement that prices lie within the unit simplex can overcome the challenge of the exponential-time convergence of price-adjustment processes in Walrasian economies, and allow for the efficient computation of Walrasian equilibrium, at least in practice.

The reader might wonder what we mean by "in practice". As it turns out the VI characterization we provide is in general discontinuous at one point in its search space, namely when the prices for all commodities are 0. As such, because it is not possible to ensure the Lipschitz-continuity or Bregman-continuity of the excess demand on the unit box in general, it is not possible to obtain polynomial-time convergence of our mirror extragradient to solve our VI without further assumptions. Nevertheless, as we discuss in the sequel, we observe the fast convergence of the mirror *extratâtonnement* process in experiments with a wide variety of competitive economies, including very large instances with Leontief consumers, for which the computation of a Walrasian equilibrium is known to be PPAD-complete (Codenotti et al., 2006; Deng and Du, 2008). These results suggest the need for a novel assumption that would explain the convergence process in practice. To this end, we introduce the pathwise Bregman-continuity assumption, a condition that requires the excess demand to be Bregman continuous along the sequence of prices generated by the mirror *extratâtonnement* process, which we show is sufficient to guarantee the polynomial-time convergence of our process.

While the pathwise Bregman-continuity assumption provides intuition on the fast convergence of the mirror *extratâtonnement* processes in practice, it is hard to verify this assumption in advance. As Bregman-continuity can be guaranteed when the price space is restricted to the unit simplex in a large class of Walrasian economies called variationally stable economies—which amongst others contain those satisfying WARP and GS—, we subsequently restrict our search space for prices to the unit simplex, and restrict our attention to competitive economies that are variationally stable on the unit simplex (i.e., those economies for which the associated VI satisfies the Minty condition) and have a **bounded elasticity of excess demand** (i.e., the percentage change in the excess demand for a percentage change in prices is bounded across all price changes). We demonstrate that under these additional assumptions, the VI is guaranteed to satisfy the Minty condition, and show that for such economies, the excess demand is Bregman continuous, thus providing the first polynomial-time convergence result for a class of price-adjustment processes in this class of Walrasian economies, which among others includes competitive economies that satisfy WGS, and more generally, WARP.

3.3.2 Variational Inequalities

Our first major contribution is introducing the class of mirror extragradient algorithms, a generalization of Korpelevich's extragradient method (Korpelevich, 1976) for solving VIs. We establish best-iterate convergence of the class of mirror extragradient algorithms to an ε-strong solution of VIs that satisfy the Minty condition and are Bregman continuous in $O(1/\varepsilon^2)$ evaluations of the optimality operator of the VI (Theorem 4.3.1). Our result generalizes the results and proof techniques of Huang and Zhang (2023) for the extragradient method, and extends the convergence results of Zhang and Dai (2023) for the unconstrained mirror extragradient method to constrained domains. In addition, to provide further justification for the convergence of the mirror extratâtonnement process in balanced economies, we establish suitable conditions for the local convergence of the mirror extragradient algorithm to an ε-strong solution of any Bregman-continuous VI that does *not* satisfy the Minty condition—to the best of our knowledge, the first result of its kind (Theorem 4.3.2).

3.3.3 Walrasian Economies

While a characterization of the set of Walrasian equilibria of any Walrasian economy as the solution set of an associated complementarity problem (i.e., a VI where the constraint set is the positive orthant) have already been known (Dafermos, 1990), for balanced economies, we provide the first computationally tractable characterization of Walrasian equilibria as the set of strong solutions of a VI that satisfies the Minty condition and whose constraint set is given by the unit box. We then apply the mirror extragradient method to obtain a novel natural price-adjustment process we call the mirror *extratâtonnement* process (Algorithm 6), and prove its convergence in all balanced economies that satisfy pathwise Bregman continuity (Corollary 5.4.1).

We then restrict our attention to a novel class of competitive economies, namely those which are variationally stable on the unit simplex, and establish the polynomial-time convergence of the mirror *extratâtonnement* process in all such economies assuming bounded elasticity of excess demand (Theorem 5.4.2). Our polynomial-time convergence result is the first such result for price-adjustment processes in the class of economies that satisfy WARP, and generalizes the well-known *tâtonnement* convergence result in competitive economies with bounded elasticity of excess demand that satisfy WGS (Codenotti et al., 2005).

We then apply the mirror *extratâtonnement* process to the Scarf economy, and prove its polynomial-time convergence to the unique Walrasian equilibrium of this economy (Corollary 5.4.4). As such, the mirror *extratâtonnement* process is the first discrete-time *natural* price-adjustment process known to converge in the Scarf economy.

Finally, we run a series of experiments on a variety of competitive economies where we verify that the pathwise Bregman-continuity assumption holds, and demonstrate that our algorithm converges to a Walrasian equilibrium at the rate predicted by our theory. Importantly, our experiments include examples of randomly initialized very large competitive economies (~ 500 consumers and ~ 500 commodities), which are known to be PPAD-complete (e.g., Leontief economies), for which we show that our algorithm finds a Walrasian equilibrium fast without failure in all cases.

3.3.4 Fisher Markets

Earlier work (Cheung, 2014; Cheung et al., 2013) has established a convergence rate of $(1 - \Theta(1))^T$ for CES Fisher markets excluding the linear and Leontief cases, and of $O\left({}^{1}\!/_{T}\right)$ for Leontief and nested[4] CES Fisher markets, where $T \in \mathbb{N}_+$ is the number of iterations for which *tâtonnement* is run. In linear Fisher markets, however, *tâtonnement* does not converge. We generalize these results by proving a convergence rate of $O({}^{(1+\epsilon^2)}\!/_{T})$, where ϵ is the maximum absolute value of the price elasticity of Hicksian demand across all buyers. Our convergence rate covers the full spectrum of homothetic Fisher markets, including mixed CES markets, i.e., CES markets with linear, Leontief, and (nested) CES buyers, unifying previously existing disparate convergence and non-convergence results. In particular, for $\epsilon = 0$, i.e., Leontief Fisher markets, we recover the best-known convergence rate of $O({}^{1}\!/_{T})$, and as $\epsilon \to \infty$, i.e., linear Fisher markets, we obtain the non-convergent behaviour of *tâtonnement* (Cole and Tao, 2019). We summarize known convergence results in light of our results in Figure 3.1a.

We observe that, in contrast to general competitive economies, in homothetic Fisher markets, concavity of the utility functions is not necessary for the existence of competitive equilibrium (Theorem 6.2.1). A computational analog of this result also holds, namely that *tâtonnement* converges in homothetic Fisher markets, even when buyers' utility functions are non-concave. Our results parallel known results on the convergence of *tâtonnement* in WGS markets, where concavity of utility functions is again not necessary for convergence (Codenotti et al., 2005).

[4] See Chapter 10 of Cheung (2014).

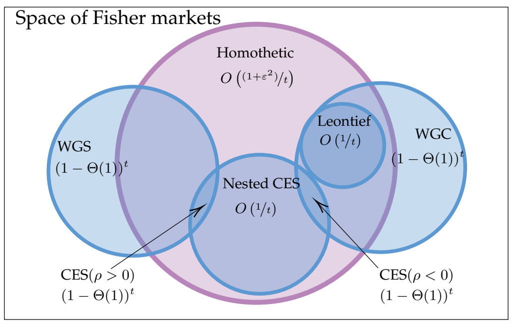

(a) The convergence rates of *tâtonnement* for different Fisher markets. We color previous contributions in blue, and our contribution in red, i.e., we study homothetic Fisher markets, where ϵ is the maximum absolute value of the price elasticity of Hicksian demand across all buyers. We note that the convergence rate for WGS markets does not apply to markets where the price elasticity of Marshallian demand is unbounded, e.g., linear Fisher markets; likewise, the convergence rate for nested CES Fisher markets does not apply to linear or Leontief Fisher markets.

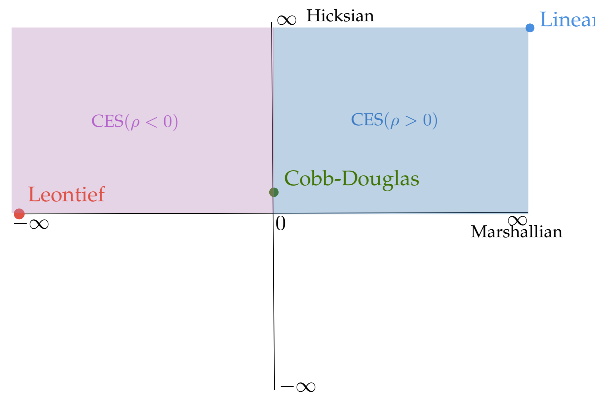

(b) Cross-price elasticity taxonomy of well-known homogeneous utility functions. There are no previously studied utility functions in the space of utility functions with negative Hicksian cross-price elasticity. Future work could investigate this space and prove faster convergence rates than those provided in this thesis. We note that our convergence result covers the entire spectrum of this taxonomy (excluding limits of the y-axis).

Figure 3.1: A summary of known results in Fisher markets.

Chapter 4

Variational Inequalities

4.1 Background

Variational inequalities (Facchinei and Pang, 2003) are a mathematical modeling framework whose study dates back to the early 1960s (Lions and Stampacchia, 1967; Hartman and Stampacchia, 1966; Browder, 1965; Grioli, 1973; Brezis and Sibony, 2011). Their utility lies in their very broad mathematical formulation which allows one to solve other mathematical modeling problems using the tools of functional analysis. They have found a great number of applications to problems in engineering and finance (Facchinei and Pang, 2003) over the years, and have recently seen an increased interest due to their novel applications in machine learning, to problems ranging from the training of generative adversarial neural networks (Goodfellow et al., 2014) to robust optimization (Ben-Tal et al., 2009).

4.1.1 Stampacchia Variational Inequality

Consider an inner product space $(\mathcal{U}, \langle \cdot, \cdot \rangle)$. A **(generalized[1]) variational inequality (VI)**, denoted $(\mathcal{X}, \mathcal{F})$, consists of a **constraint set** $\mathcal{X} \subseteq \mathcal{U}$ and an **optimality operator** $\mathcal{F} : \mathcal{U} \rightrightarrows \mathcal{U}^*$. For notational convenience, for any $\boldsymbol{x} \in \mathcal{X}$, we denote any arbitrary element of $\mathcal{F}(\boldsymbol{x})$ by $\boldsymbol{f}(\boldsymbol{x})$, and denote the variational inequality by $(\mathcal{X}, \boldsymbol{f})$ if $\mathcal{F}$ is singleton-valued.

[1]When $\mathcal{F}$ is singleton-valued a generalized variational inequality is simply called a variational inequality. As our computational results will be limited to generalized variational inequalities where $\mathcal{F}$ is singleton-valued, for simplicity, we will refer to generalized variational inequalities simply as variational inequalities.

Any VI $(\mathcal{X}, \mathcal{F})$ defines a problem known as the **(generalized) Stampacchia variational inequality (SVI)** (Lions and Stampacchia, 1967):

$$\text{Find } \boldsymbol{x}^* \in \mathcal{X} \text{ and } \boldsymbol{f}(\boldsymbol{x}^*) \in \mathcal{F}(\boldsymbol{x}^*) \text{ s.t. } \quad \text{for all } \boldsymbol{x} \in \mathcal{X}, \qquad \langle \boldsymbol{f}(\boldsymbol{x}^*), \boldsymbol{x} - \boldsymbol{x}^* \rangle \geq 0 \tag{4.1}$$

A solution to a SVI is called a **strong solution** of the variational inequality $(\mathcal{X}, \mathcal{F})$. Just like in convex optimization settings (see Section 1.1.2 of Nesterov (1998)), in practice, it is not possible to compute an exact strong solution to a VI $(\mathcal{X}, \mathcal{F})$ in finite-time, and as such we have to resort to approximate solutions which we call the ε-strong solution. Note that in the following definition, in line with the literature (see, for instance Section 1.2 of Diakonikolas (2020)), the inequality is negated (and as such inverted).

Definition 4.1.1 [Strong Solution]**.**
Given an **approximation parameter** $\varepsilon \geq 0$, an ε-**strong** (or **Stampacchia**) **solution** of the VI $(\mathcal{X}, \mathcal{F})$ is an $\boldsymbol{x}^* \in \mathcal{X}$ for which there exists $\boldsymbol{f}(\boldsymbol{x}^*) \in \mathcal{F}(\boldsymbol{x}^*)$ s.t. $\max_{\boldsymbol{x} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}^*), \boldsymbol{x}^* - \boldsymbol{x} \rangle \leq \varepsilon$.

A 0-strong solution is simply called a **strong solution**. We denote the set of ε-strong (respectively, the set of strong) solutions of a VI $(\mathcal{X}, \mathcal{F})$ by $\mathcal{SVI}_\varepsilon(\mathcal{X}, \mathcal{F})$ (respectively, $\mathcal{SVI}(\mathcal{X}, \mathcal{F})$).

To understand this definition, first take a look at the following equivalent formulations:

$$\begin{aligned}
&\max_{\boldsymbol{x} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}^*), \boldsymbol{x}^* - \boldsymbol{x} \rangle \leq \varepsilon \\
\iff &\max_{\boldsymbol{x} \in \mathcal{X}} - \langle \boldsymbol{f}(\boldsymbol{x}^*), \boldsymbol{x} - \boldsymbol{x}^* \rangle \leq \varepsilon \\
\iff &- \min_{\boldsymbol{x} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}^*), \boldsymbol{x} - \boldsymbol{x}^* \rangle \leq \varepsilon \\
\iff &\min_{\boldsymbol{x} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}^*), \boldsymbol{x} - \boldsymbol{x}^* \rangle \geq -\varepsilon
\end{aligned}$$

The final inequality tells us that any first-order deviations from $\boldsymbol{x}^*$ can change the objective by at most $-\varepsilon$ (i.e., decrease the objective at most by ε). Since the LHS is negative (simply plug in $\boldsymbol{x} = \boldsymbol{x}^*$), when $\varepsilon = 0$, first-order deviations away from a solution cannot decrease the objective further.

A large number of mathematical optimization problems can be cast as VI problems, and as such they have found a large number of applications. We will explore a number of these applications to game theory in this thesis, and mention now only a simple application to convex optimization for illustrative purposes.

Example 4.1.1 [Convex Optimization as a VI]**.**

Consider a convex optimization problem $(\mathcal{X}, h)$, where $\mathcal{X} \subseteq \mathcal{U}$ is a non-empty, compact, and convex constraint set, and $h : \mathcal{U} \to \mathbb{R}$ is a continuous and convex objective function: i.e.,

$$\min_{\boldsymbol{x} \in \mathcal{X}} h(\boldsymbol{x})$$

Any ε-minimum $\boldsymbol{x}^* \in \mathcal{X}$ of $(\mathcal{X}, h)$ s.t. $h(\boldsymbol{x}^*) - \min_{\boldsymbol{x} \in \mathcal{X}} h(\boldsymbol{x}) \leq \varepsilon$ satisfies the following necessary and sufficient optimality conditions (see, for instance Section 2 of Crespi et al. (2005)):

$$\langle \partial h(\boldsymbol{x}^*), \boldsymbol{x}^* - \boldsymbol{x} \rangle \leq \varepsilon \qquad\qquad \forall \boldsymbol{x} \in \mathcal{X}, \exists \partial h(\boldsymbol{x}^*) \in \mathcal{D}h(\boldsymbol{x}^*)$$

Taking a maximum over $\boldsymbol{x} \in \mathcal{X}$, we can then see that the set of ε-minima of $(\mathcal{X}, h)$ corresponds to the set of ε-strong solutions $\mathcal{SVI}_\varepsilon(\mathcal{X}, \mathcal{D}h)$ of the VI $(\mathcal{X}, \mathcal{D}h)$, where the optimality operator is given by the subdifferential correspondence of the objective h.

Going further, if $(\mathcal{X}, h)$ is not a convex optimization problem, but is instead a weakly-convex optimization problem, then the optimality conditions are only sufficient, in which case the set of ε-strong solutions $\mathcal{SVI}_\varepsilon(\mathcal{X}, \mathcal{D}h)$ of the VI $(\mathcal{X}, \mathcal{D}h)$ is called the set of ε-**stationary points** of $(\mathcal{X}, h)$.

Strong solutions can be shown to exist in a broad of class of VIs known as continuous VIs.

Definition 4.1.2 [Continuous VIs]**.**

A **continuous** VI is a VI $(\mathcal{X}, \mathcal{F})$ such that:

1. $\mathcal{X}$ is non-empty, compact, and convex
2. $\mathcal{F}$ is upper hemicontinuous, non-empty-valued, compact-valued, and convex-valued

The proof of existence of a strong solution in continuous VIs relies on a fixed-point argument applied to a mapping whose fixed points correspond to strong solutions of the VI, which can be shown to exist by the Glicksberg-Kakutani fixed point theorem (see Theorem 2.4.1). We refer the reader to Theorem 2.2.1 of Facchinei and Pang (2003) for details.

Theorem 4.1.1 [Existence of Strong Solution (Theorem 2.2.1 of Facchinei and Pang (2003))]**.** Every continuous VI $(\mathcal{X}, \mathcal{F})$ has at least one strong solution, i.e., $\mathcal{SVI}(\mathcal{X}, \mathcal{F}) \neq \emptyset$.

4.1.2 Minty Variational Inequality

An alternative but related problem formulation for VIs is the **(generalized) Minty variational inequality (MVI)** (Minty, 1967). Given a VI $(\mathcal{X}, \mathcal{F})$, the MVI is defined as:

$$\text{Find } \boldsymbol{x}^* \in \mathcal{X} \text{ s.t. for all } \boldsymbol{x} \in \mathcal{X} \text{ there exists } \boldsymbol{f}(\boldsymbol{x}) \in \mathcal{F}(\boldsymbol{x}) \text{ s.t. } \langle \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x}^* - \boldsymbol{x} \rangle \leq 0 \tag{4.2}$$

A solution to a MVI is called the weak solution, for which we similarly can define an approximate variant for computational purposes.

Definition 4.1.3 [Weak (or Minty) Solution]**.**
Given a VI $(\mathcal{X}, \mathcal{F})$ and an **approximation parameter** $\varepsilon \geq 0$, an ε**-weak (or Minty) solution** is an $\boldsymbol{x}^* \in \mathcal{X}$ for which there exists $\boldsymbol{f}(\boldsymbol{x}) \in \mathcal{F}(\boldsymbol{x})$ s.t. $\max_{\boldsymbol{x} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x}^* - \boldsymbol{x} \rangle \leq \varepsilon$.

A 0-weak solution to the VI is simply called a **weak solution**. We denote the set of ε-weak (respectively, the set of weak) solutions a VI $(\mathcal{X}, \mathcal{F})$ by $\mathcal{MVI}_\varepsilon(\mathcal{X}, \mathcal{F})$ (respectively, $\mathcal{MVI}(\mathcal{X}, \mathcal{F})$).

The inequality $\max_{\boldsymbol{x} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x}^* - \boldsymbol{x} \rangle \leq \varepsilon$ says that any movement from any $\boldsymbol{x}$ to an $\boldsymbol{x}^*$ can increase the objective by at most ε. Now since the LHS is non-negative (simply plug in $\boldsymbol{x} = \boldsymbol{x}^*$), when $\varepsilon = 0$, any movement from another variable to a solution cannot increase the objective; it can only decrease it.

In continuous VIs (see Definition 4.1.2), the set of weak solutions is a subset of set of strong solutions, i.e., the MVI is a refinement of the SVI. However, we note that a weak solution is in general not guaranteed to exist in continuous VIs.

Remark 4.1.1 [Strong vs. Weak Solutions]**.**

It might seem like be a misnomer that the solutions of the MVI are called weak solutions as they are a subset of the strong solutions; however, beyond finite dimensional settings (e.g., the Euclidean space setting of interest here), the set of weak and strong solutions may be totally unrelated, as they are not guaranteed to have a non-empty intersection. As a VI is a mathematical framework for modeling other mathematical problems using the tools of functional analysis, a weak solution should be interpreted as "weak" from a modeling perspective, in the following sense.

In a VI $(\mathcal{X}, \mathcal{F})$, we are given an optimality operator $\mathcal{F}$ which maps an $\boldsymbol{x} \in \mathcal{X}$ to a set of optimality conditions. These optimality conditions are then tested against other $\boldsymbol{x}' \in \mathcal{X}$ using an inner product. For strong solutions of the VI, the candidate solution we are seeking, say $\boldsymbol{x}^*$, is mapped by the optimality operator $\mathcal{F}$; as such, the operator must be well defined at $\boldsymbol{x}^*$, and more importantly, the optimality conditions $\boldsymbol{f}(\boldsymbol{x}^*)$ must be satisfied by the candidate solution $\boldsymbol{x}^*$. In contrast, for weak solutions of the VI, the optimality conditions $\boldsymbol{f}(\boldsymbol{x}^*)$ need not be satisfied by $\boldsymbol{x}^*$, and the optimality operator need not even be defined (i.e., $\mathcal{F}(\boldsymbol{x}^*)$ might be empty).

That is, a weak solution is weak as there is no notion of optimality that can be attributed to it without testing it against other $\boldsymbol{x}' \in \mathcal{X}$, while a strong solution is itself strong as its optimality can almost be determined independently of other $\boldsymbol{x}' \in \mathcal{X}$: the optimality conditions fully characterize a solution. As such, when we cast a mathematical problem as a VI, the set of weak solutions might be underspecified, and thus nothing is to be gained by casting the original problem as a VI. To see this concretely, recall Example 4.1.1, and suppose that the objective of the optimization problem is non-convex. The set of weak solutions to the corresponding VI corresponds to the set of global solutions of the optimization problem (see, for instance, Proposition 2.2 of Crespi et al. (2005)), while the set of strong solutions corresponds to its stationary points. The MVI formulation of the optimization problem is thus no more informative than the original problem. In contrast, the stationary points of the optimization problem, which correspond to the strong solutions of the VI provide additional information, namely the necessary conditions that must be satisfied by the solutions of the original optimization problem. Perhaps, more importantly, the optimality of a local solution to the optimization problem (and as such of the associated SVI problem) can be determined almost independently of other points: the first-order conditions fully characterize a local solution. In contrast, the optimality of a global solution to the optimization problem (and as such of the associated MVI) requires one to test a candidate solution against *all* other points. For additional discussion, we refer the reader to Zang and Avriel (1975).

4.1.3 Generalized Monotonicity Properties of Variational Inequalities

The following additional properties of VIs define important properties of the set of strong and weak solutions of VIs, and will be relevant in the sequel. In particular, a large number of VIs satisfy a number of monotonicity conditions which makes them more analytically tractable.

Definition 4.1.4 [Monotone, Pseudomonotone and Quasimonotone VIs]**.**
A VI $(\mathcal{X}, \mathcal{F})$ is **monotone** (respectively, **pseudomonotone** / **quasimonotone**) iff the optimality operator $\mathcal{F}$ is monotone (respectively, **pseudomonotone** / **quasimonotone**).[2]

Another common and more general property for the analysis of VIs, known as the Minty condition, is simply the existence of a Minty solution.

Definition 4.1.5 [Minty's Condition]**.**
A VI $(\mathcal{X}, \mathcal{F})$ satisfies the **Minty condition** iff $\mathcal{MVI}(\mathcal{X}, \mathcal{F}) \neq \emptyset$.

Example 4.1.2 [Where Minty's condition does not hold]**.**
Consider the VI $(\mathcal{X}, f)$ associated with the concave minimization problem $\min_{x\in[-1,1]} -x^2$, with $\mathcal{X} = [-1, 1]$ and $f(x) = -2x$. Solutions to this minimization problem occur at $x^* = -1$ and $x^* = 1$. Now, note that for $x^* = -1$ and $x \in (0, 1]$, we have $\langle f(x), x - x^*\rangle < 0$. Similarly,

[2]For a definition of monotone, pseudomonotone and quasimonotone operators, Section 2.8.

for $x^* = 1$ and $x \in [-1, 0)$, we have $\langle f(x), x - x^* \rangle < 0$. Since any Minty solution is a (global) solution of the minimization problem, a Minty solution cannot exist.

With these definitions in order, we summarize the following known properties of the solution sets of VIs.

Remark 4.1.2 [Solution Set Properties]**.**
Let $\varepsilon \geq 0$, then the following implications hold:

- $(\mathcal{X}, \mathcal{F})$ is continuous $\implies \mathcal{SVI}(\mathcal{X}, \mathcal{F}) \neq \emptyset$ (Theorem 2.2.1 of Facchinei and Pang (2003)))
- $(\mathcal{X}, \mathcal{F})$ is continuous $\implies \mathcal{MVI}(\mathcal{X}, \mathcal{F}) \subseteq \mathcal{SVI}(\mathcal{X}, \mathcal{F})$
- $(\mathcal{X}, \mathcal{F})$ is monotone $\implies \mathcal{SVI}_\varepsilon(\mathcal{X}, \mathcal{F}) \subseteq \mathcal{MVI}_\varepsilon(\mathcal{X}, \mathcal{F})$
- $(\mathcal{X}, \mathcal{F})$ is pseudomonotone $\implies \mathcal{SVI}(\mathcal{X}, \mathcal{F}) \subseteq \mathcal{MVI}(\mathcal{X}, \mathcal{F})$
- $(\mathcal{X}, \mathcal{F})$ is quasimonotone with $\mathcal{X}$ non-empty and compact $\implies$ Minty condition (Lemma 3.1 of (He, 2017))
- If $\mathcal{SVI}(\mathcal{X}, \mathcal{F}) \neq \emptyset$, then monotone $\implies$ pseudomonotone $\implies$ Minty's condition

Note that while it has become common place to use the Minty condition in the analysis of VIs as it is very general (see, for instance, He et al. (2022)), in all the applications discussed in this thesis, the Minty condition can be replaced by the assumption that the VI $(\mathcal{X}, \mathcal{F})$ is quasimonotone with $\mathcal{X}$ non-empty, and compact, by Lemma 3.1 and Proposition 3.1 of (He, 2017).

4.2 Algorithms for Variational Inequalities

We now turn our attention to the computation of solutions to variational inequalities. In what follows, for simplicity, we will restrict ourselves to VIs $(\mathcal{X}, \mathcal{F})$ in which $\mathcal{F}$ is singleton-valued, which we will denote simply as $(\mathcal{X}, \boldsymbol{f})$. In future work, the algorithms and results provided in this chapter could be extended to the more general non-singleton-valued VI setting.

4.2.1 Computational Model

In this thesis, we will consider two classes of methods to solve VIs, first-order and second-order methods which both belong to the class of kth-order methods.

Definition 4.2.1 [kth-order methods]**.**

Given some $k \in \mathbb{N}_{++}$, a VI $(\mathcal{X}, \mathcal{F})$ for which the derivatives $\{\nabla^j \boldsymbol{f}\}_{j=1}^{k-1}$ are well defined, and an initial iterate $\boldsymbol{x}^{(0)} \in \mathcal{X}$, a k**th-order method** $\boldsymbol{\mu}$ consists of an update function which generates the sequence of iterates $\{\boldsymbol{x}^{(t)}\}_t$ given by: for all $t = 0, 1, \ldots$,

$$\boldsymbol{x}^{(t+1)} \doteq \boldsymbol{\mu}\left(\bigcup_{i=0}^{t} (\boldsymbol{x}^{(i)}, \{\nabla^j \boldsymbol{f}(\boldsymbol{x}^{(i)})\}_{j=0}^{k-1})\right)$$

The computational complexity results in this chapter will rely on the following computational model which has been broadly adopted in the literature (see, for instance, Cai et al. (2022)).

Definition 4.2.2 [VI Computational Model]**.**

Given a VI $(\mathcal{X}, \boldsymbol{f})$ and a kth-order method $\boldsymbol{\mu}$, the computational complexity of $(\mathcal{X}, \boldsymbol{f})$ is measured in term of the number of evaluations of the the functions $\boldsymbol{f}, \nabla \boldsymbol{f}, \ldots, \nabla^k \boldsymbol{f}$.

Remark 4.2.1.

In line with the literature, the computational model we consider thus assumes that any other operation such as a (Bregman) projection onto a set is a constant cost operation.

The computational results in the literature, as well as the results we will present in this chapter, hold for the following class of VIs.

Definition 4.2.3 [Lipschitz-Continuous VIs]**.**

Given a modulus of continuity $\lambda \geq 0$, a λ-**Lipschitz-continuous** VI is a VI $(\mathcal{X}, \boldsymbol{f})$ such that:

1. $\mathcal{X}$ is non-empty, compact, and convex
2. $\boldsymbol{f}$ is λ-Lipschitz continuous

4.2.2 Related Work

Historically, the goal of the literature on solution methods for VIs has been to devise algorithms which are asymptotically guaranteed to converge to a strong or weak solution (Brezis and Brezis, 2011). An overwhelming majority of these works have focused on first-order methods for computing solutions of VIs, with higher-order methods having been considered only in recent years (see, for instance, He et al. (2022); Huang and Zhang (2022)). While a strong solution of a VI is guaranteed to exist in continuous VIs, most results on the computational complexity of strong solutions concern the class of monotone VIs (see, for instance Cai et al. (2022)) with a few works focusing on VIs that satisfy the Minty condition (see, for instance, Diakonikolas (2020)).

The canonical algorithm to solve VIs is the projected gradient method (Cauchy et al., 1847; Nesterov, 1998) (also known under the names of the subgradient method, gradient descent ascent, or the Arrow-Hurwicz-Uzawa method (Arrow and Hurwicz, 1958; Arrow et al., 1958)). While asymptotic convergence of the projected gradient method can be shown for a subset of monotone VIs known as strongly monotone VIs, in general monotone VIs, only ergodic asymptotic convergence (i.e., asymptotic convergence of the averaged iterates) to a solution[3] can be guaranteed. The earliest known algorithm with asymptotic convergence guarantees to a solution of a monotone VI is the extragradient method, attributed to Korpelevich (1976). Following this early success, Popov (1980) introduced a closely related algorithm called the optimistic gradient method, which he also showed to converge to a solution. These initial extragradient and optimistic gradient algorithms would eventually become much more sophisticated, with a large body of work appearing on asymptotic convergence guarantees for variants of these earlier methods (e.g., (Solodov and Svaiter, 1999)).

More recently, the literature has turned its attention to algorithms with non-asymptotic guarantees, and in particular to ones that are guaranteed to compute an ε-strong or ε-weak

[3]Recall that for monotone VIs, the set of strong and weak solutions are equal: as such, "solution" here refers to both strong and weak solutions.

solution of a VI in polynomial-time, i.e., in a number of evaluations of the optimality operator $\mathcal{F}$ which is polynomial in the inverse of the approximation parameter $1/\varepsilon$, the dimensionality n of the constraint set, and other relevant assumption-specific parameters, such as an upper bound on all of the values of the optimality operator. One of the earliest results in this direction was given by Nemirovski (2004), who introduced the conceptual Mirror-Prox method, an elegant generalization of the Extragradient Method, and established that ε-strong and ε-weak solutions can be computed in $O(1/\varepsilon)$ operations by averaging the iterates of the algorithm under the assumptions that the VI is monotone and the optimality operator is Lipschitz-continuous. Nemirovski's work inspired a large body of work on more sophisticated algorithms (e.g., (Auslender and Teboulle, 2005; Diakonikolas, 2020)) for monotone VIs, and better computational results for the projection method (Gidel et al., 2018), the extragradient method (Gorbunov et al., 2022a; Golowich et al., 2020a; Cai et al., 2022) and the optimistic gradient method (Gorbunov et al., 2022b).

A number of works have considered first-order methods to compute a strong solution (e.g., Loizou et al. (2021); He et al. (2022); Diakonikolas (2020)), or a stationary point, of the VI[4] (e.g., Liu et al. (2021)) that also satisfies the Minty condition. The first-order methods considered by this more recent line of work on non-monotone variational inequalities include the extragradient method (e.g., (Wang and Ma, 2024; Ofem et al., 2023)), Tseng's method (e.g., (Censor et al., 2011; Thong et al., 2020; Uzor et al., 2023; Dung et al., 2024; Aremu et al., 2024)), and the optimistic gradient method (e.g., (Lin and Jordan, 2022)) and its variants.

[4]A (ε, δ) stationary point of a VI $(\mathcal{X}, \mathcal{F})$ is a point $\boldsymbol{x}^* \in \mathcal{X}$ s.t. for some $\delta \geq 0$, there exists $\boldsymbol{x} \in \mathcal{B}_\delta(\boldsymbol{x}^*)$ and $\boldsymbol{x}$ is an ε-strong solution. Convergence to this weaker solution concept is necessary for VIs in which $\mathcal{F}$ is not singleton-valued for technical reasons, and any future work that seeks to generalize the results in this section should adopt this weaker definition to prove convergence results.

4.3 First-Order Methods

We will at present focus on first-order methods for VIs.

4.3.1 Mirror Gradient Algorithm

The canonical class of first-order methods that can be used to compute strong solutions for VIs is the class of **mirror gradient algorithms** (Algorithm 1, (Nemirovskij and Yudin, 1983)). These methods are parameterized by a kernel function $h : \mathcal{X} \to \mathbb{R}$, which induces a a Bregman divergence $\mathrm{div}_h : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ that defines the algorithm's update function $\boldsymbol{\mu}$.

Definition 4.3.1 [Bregman Divergence]**.**
Given a set $\mathcal{X}$ and a **kernel function** $h : \mathcal{X} \to \mathbb{R}$, the **Bregman divergence** $\mathrm{div}_h : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ associated with h is defined as:

$$\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{y}) \doteq h(\boldsymbol{x}) - h(\boldsymbol{y}) - \langle \nabla h(\boldsymbol{y}), \boldsymbol{x} - \boldsymbol{y} \rangle$$

We note the following properties of the Bregman divergence. For additional background, see, for instance Zhang and He (2018).

Remark 4.3.1 [Properties of the Bregman divergence]**.**
When h is convex, for all $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{X}$, the Bregman divergence is positive, i.e., $\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{y}) \geq 0$. Further, if h is strictly convex, then $\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{y}) = 0$ iff $\boldsymbol{x} = \boldsymbol{y}$. In addition, if h is μ-strongly convex, then for all $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{X}$, we have $\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{y}) \geq \frac{\mu}{2} \|\boldsymbol{x} - \boldsymbol{y}\|^2$.

Algorithm 1 Mirror Gradient Algorithm

Input: $\mathcal{X}, \boldsymbol{f}, h, \tau, \eta, \boldsymbol{x}^{(0)}$

Output: $\{\boldsymbol{x}^{(t)}\}_t$

1: **for** $t = 1, \dots, \tau$ **do**
2: $\quad \boldsymbol{x}^{(t+1)} \leftarrow \arg\min_{\boldsymbol{x} \in \mathcal{X}} \left\{ \left\langle \boldsymbol{f}(\boldsymbol{x}^{(t)}), \boldsymbol{x} - \boldsymbol{x}^{(t)} \right\rangle + \frac{1}{2\eta} \mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(t)}) \right\}$
return $\{\boldsymbol{x}^{(t)}\}_t$

When the kernel function is chosen s.t. $h(\boldsymbol{x}) \doteq \frac{1}{2} \|\boldsymbol{x}\|^2$, the Bregman divergence corresponds to the Euclidean square norm, i.e., $\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{y}) \doteq \|\boldsymbol{x} - \boldsymbol{y}\|^2$, in which case the mirror gradient method reduces to the well-known **projected gradient** method (Algorithm 2, (Cauchy et al., 1847)).

Algorithm 2 Project Gradient Algorithm

Input: $\mathcal{X}, \boldsymbol{f}, \tau, \eta, \boldsymbol{x}^{(0)}$

Output: $\{\boldsymbol{x}^{(t)}\}_t$

1: Initialize $\boldsymbol{x}^{(1)} \in \mathcal{X}$ arbitrarily

2: **for** $t = 1, \ldots, \tau$ **do**

3: $\quad \boldsymbol{x}^{(t+1)} \leftarrow \Pi_{\mathcal{X}}\left[\boldsymbol{x}^{(t)} - \eta \boldsymbol{f}(\boldsymbol{x}^{(t)})\right] = \arg\min_{\boldsymbol{x} \in \mathcal{X}} \left\{ \left\langle \boldsymbol{f}(\boldsymbol{x}^{(t)}), \boldsymbol{x} - \boldsymbol{x}^{(t)} \right\rangle + \frac{1}{2\eta} \|\boldsymbol{x} - \boldsymbol{x}^{(t)}\|^2 \right\}$

return $\{\boldsymbol{x}^{(t)}\}_t$

Unfortunately, while the average of the iterates of the mirror gradient method can be shown to converge asymptotically to a strong solution for monotone VIs with a Lipschitz-continuous optimality operator, it is only possible to prove polynomial-time computation of an ε-weak solution in such VIs, which does not necessarily imply convergence to an ε-strong solution (see, for instance, Proposition 8 and Appendix D of Liu et al. (2021)). More importantly, in general the sequence of iterates generated by the mirror gradient method is not guaranteed to converge, as shown by the following example.

Example 4.3.1 [Non-Convergence of Mirror Gradient Method]**.**
Consider the VI $(\mathcal{X}, \boldsymbol{f})$ with $\mathcal{X} \doteq \mathbb{R}^2$ and $\boldsymbol{f}(x, y) = (-y, x)$. For this VI, we have $\mathcal{SVI}(\mathcal{X}, \boldsymbol{f}) = \mathcal{MVI}(\mathcal{X}, \boldsymbol{f}) = \{(0,0)\}$. Suppose that $(\boldsymbol{x}^{(0)}, \boldsymbol{y}^{(0)}) \neq (0,0)$, then for any $\eta > 0$, the iterates generated by the gradient method are given by:

$$(x^{(t)}, y^{(t)}) \doteq \left(x^{(0)} - \eta \sum_{k=1}^{t} y^{(k-1)}, y^{(0)} + \eta \sum_{k=1}^{t} x^{(k-1)} \right) \qquad \forall t \in \mathbb{N}_{++} \tag{4.3}$$

and as such are unbounded, i.e., $\|(\boldsymbol{x}^{(t)}, \boldsymbol{y}^{(t)})\| \to \infty$.

4.3.2 Mirror Extragradient Algorithm

As the iterates of the mirror gradient method do not asymptotically converge to a strong or weak solution, and it is not possible to obtain polynomial-time computation of an ε-strong solution by averaging the iterates, we now introduce a novel class of first order methods, namely the class of **mirror extragradient algorithms**, which similar to the class of mirror gradient methods, are parameterized by a kernel function $h : \mathcal{X} \to \mathbb{R}$ that induces a Bregman divergence $\text{div}_h : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ defining the update function $\boldsymbol{\mu}$ of the algorithm.

Algorithm 3 Mirror Extragradient Algorithm

Input: $\mathcal{X}, \boldsymbol{f}, h, \tau, \eta, \boldsymbol{x}^{(0)}$

Output: $\{\boldsymbol{x}^{(t+0.5)}, \boldsymbol{x}^{(t+1)}\}_t$

1: **for** $t = 1, \ldots, \tau$ **do**

2: $\quad \boldsymbol{x}^{(t+0.5)} \leftarrow \arg\min_{\boldsymbol{x} \in \mathcal{X}} \left\{ \left\langle \boldsymbol{f}(\boldsymbol{x}^{(t)}), \boldsymbol{x} - \boldsymbol{x}^{(t)} \right\rangle + \frac{1}{2\eta} \text{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(t)}) \right\}$

3: $\quad \boldsymbol{x}^{(t+1)} \leftarrow \arg\min_{\boldsymbol{x} \in \mathcal{X}} \left\{ \left\langle \boldsymbol{f}(\boldsymbol{x}^{(t+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(t)} \right\rangle + \frac{1}{2\eta} \text{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(t)}) \right\}$

return $\{\boldsymbol{x}^{(t+0.5)}, \boldsymbol{x}^{(t+1)}\}_t$

The mirror extragradient algorithm (Algorithm 3) generalizes the well-known extragradient algorithm which is known to asymptotically converge to a strong solution (Popov, 1980), and allows for the polynomial-time computation of an ε-strong solution (Nemirovski, 2004; Golowich et al., 2020b; Cai et al., 2022). In particular, when the kernel function for the mirror extragradient method is chosen s.t. $h(\boldsymbol{x}) \doteq \frac{1}{2} \|\boldsymbol{x}\|^2$, the Bregman divergence corresponds to the Euclidean square norm, i.e., $\text{div}_h(\boldsymbol{x}, \boldsymbol{y}) \doteq \|\boldsymbol{x} - \boldsymbol{y}\|^2$, in which case the mirror extragradient method reduces to the extragradient method (Algorithm 4).

Algorithm 4 (Projected) Extragradient Algorithm

Input: $\mathcal{X}, \boldsymbol{f}, \tau, \eta, \boldsymbol{x}^{(0)}$

Output: $\{\boldsymbol{x}^{(t+0.5)}, \boldsymbol{x}^{(t+1)}\}_t$

1: **for** $t = 1, \ldots, \tau$ **do**

2: $\quad \boldsymbol{x}^{(t+0.5)} \leftarrow \Pi_{\mathcal{X}} \left[\boldsymbol{x}^{(t)} - \eta \boldsymbol{f}(\boldsymbol{x}^{(t)}) \right] = \arg\min_{\boldsymbol{x} \in \mathcal{X}} \left\{ \left\langle \boldsymbol{f}(\boldsymbol{x}^{(t)}), \boldsymbol{x} - \boldsymbol{x}^{(t)} \right\rangle + \frac{1}{2\eta} \|\boldsymbol{x} - \boldsymbol{x}^{(t)}\|^2 \right\}$

3: $\quad \boldsymbol{x}^{(t+1)} \leftarrow \Pi_{\mathcal{X}} \left[\boldsymbol{x}^{(t)} - \eta \boldsymbol{f}(\boldsymbol{x}^{(t+0.5)}) \right] = \arg\min_{\boldsymbol{x} \in \mathcal{X}} \left\{ \left\langle \boldsymbol{f}(\boldsymbol{x}^{(t+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(t)} \right\rangle + \frac{1}{2\eta} \|\boldsymbol{x} - \boldsymbol{x}^{(t)}\|^2 \right\}$

return $\{\boldsymbol{x}^{(t+0.5)}, \boldsymbol{x}^{(t+1)}\}_t$

A seminal result by Nemirovski (2004) shows that the average of the iterates output by the extragradient algorithm are an ε-strong solution for any monotone VI with a Lipschitz-continuous optimality operator when the algorithm is run for $\tau \in O(1/\varepsilon)$ time-steps. Additionally, Golowich et al. (2020b); Cai et al. (2022) show that in the same setting, the tth iterate is an ε-strong solution when the algorithm is run for $\tau \in O(1/\varepsilon^2)$ time-steps. More recently, Huang and Zhang (2023) have extended these polynomial-time computation result to VIs which satisfy the weaker Minty condition rather than monotonicity assumption, showing that there exists some $k \leq \tau$ s.t. the kth iterate of the extragradient algorithm is an ε-strong solution for $\tau \in O(1/\varepsilon^2)$.[5] We extend at present this result to the mirror extragradient method. In order to prove, our result we prove several lemmas. We start with a technical lemma on Bregman divergences.

Lemma 4.3.1 [Bregman Triangle Lemma]**.**
Consider the Bregman divergence $\mathrm{div}_h : \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ associated with a differentiable kernel function $h : \mathcal{X} \to \mathbb{R}$. For all $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \in \mathcal{X}$,

$$\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{z}) + \mathrm{div}_h(\boldsymbol{y}, \boldsymbol{x}) - \mathrm{div}_h(\boldsymbol{y}, \boldsymbol{z}) = \langle \nabla h(\boldsymbol{x}) - \nabla h(\boldsymbol{z}), \boldsymbol{x} - \boldsymbol{y} \rangle. \tag{4.4}$$

Proof of Lemma 4.3.1

For all $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \in \mathcal{X}$, we have:

$$\begin{aligned}
&\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{z}) + \mathrm{div}_h(\boldsymbol{y}, \boldsymbol{x}) - \mathrm{div}_h(\boldsymbol{y}, \boldsymbol{z}) \\
&= [h(\boldsymbol{x}) - h(\boldsymbol{z}) - \langle \nabla h(\boldsymbol{z}), \boldsymbol{x} - \boldsymbol{z} \rangle] + [h(\boldsymbol{y}) - h(\boldsymbol{x}) - \langle \nabla h(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle] \\
&- [h(\boldsymbol{y}) - h(\boldsymbol{z}) - \langle \nabla h(\boldsymbol{z}), \boldsymbol{y} - \boldsymbol{z} \rangle] \\
&= -\langle \nabla h(\boldsymbol{z}), \boldsymbol{x} - \boldsymbol{z} \rangle - \langle \nabla h(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle + \langle \nabla h(\boldsymbol{z}), \boldsymbol{y} - \boldsymbol{z} \rangle \\
&= \langle \nabla h(\boldsymbol{z}) - \nabla h(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle
\end{aligned}$$

With the above technical lemma in hand, we are now ready to prove a progress lemma for the mirror extragradient method, which describes how the algorithm progresses from one iteration to another. Note that under the Minty condition, the following lemma

[5] This type of convergence is known as a best-iterate convergence.

implies convergence to a weak solution, since setting $\boldsymbol{x} \doteq \boldsymbol{x}^* \in \mathcal{MVI}(\mathcal{X}, \boldsymbol{f})$, we obtain $\mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k)}) > \mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k+1)})$, for all $k \in \mathbb{N}$ (i.e., the distance to the weak solution $\boldsymbol{x}^*$ is strictly decreasing). Below, we first introduce a condition necessary for the result of the lemma to hold, and then we state the lemma.

Definition 4.3.2 [Pathwise Bregman-Continuity]**.**
A VI $(\mathcal{X}, \boldsymbol{f})$ is **pathwise Bregman continuous** over the outputs of the mirror extragradient method $\{\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}\}_t$, i.e., there exists $\lambda \geq 0$, s.t. for all $k \in \mathbb{N}$, $\frac{1}{2}\|\boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}) - \boldsymbol{f}(\boldsymbol{x}^{(k)})\|^2 \leq \lambda^2 \,\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})$.

Lemma 4.3.2 [Mirror Extragradient Progress]**.**
Consider the mirror extragradient algorithm (Algorithm 3) run on a VI $(\mathcal{X}, \boldsymbol{f})$ with a 1-strongly-convex kernel function h, a step size $\eta > 0$, and a time horizon $\tau \in \mathbb{N}$. If the outputs $\{\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k+1)}\}_k$ are such that the $(\mathcal{X}, \boldsymbol{f})$ is pathwise Bregman-continuous, i.e., there exists $\lambda \geq 0$, s.t. for all $k \in \mathbb{N}$ $\frac{1}{2}\|\boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}) - \boldsymbol{f}(\boldsymbol{x}^{(k)})\|^2 \leq \lambda^2 \,\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})$, then for all $k \in \mathbb{N}$ and $\boldsymbol{x} \in \mathcal{X}$, the following inequality holds for all $k = 0, 1, \ldots$:

$$\begin{aligned}
&\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k)}) - \mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k+1)}) \\
&\geq \eta\langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}\rangle + \left(1 - (\eta\lambda)^2\right) \mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})
\end{aligned}$$

Proof of Lemma 4.3.2

By the first-order optimality condition at $\boldsymbol{x}^{(k+0.5)}$, for all $\boldsymbol{x} \in \mathcal{X}$,

$$\langle \boldsymbol{f}(\boldsymbol{x}^{(k)}) + \frac{1}{\eta}\nabla h(\boldsymbol{x}^{(k+0.5)}) - \nabla h(\boldsymbol{x}^{(k)}), \boldsymbol{x} - \boldsymbol{x}^{(k+0.5)}\rangle \geq 0 \enspace .$$

Substituting $\boldsymbol{x} = \boldsymbol{x}^{(k+1)}$ above and rearranging terms yields:

$$\begin{aligned}
&\langle \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k+0.5)}\rangle && (4.5)\\
&\geq \frac{1}{\eta}\langle \nabla h(\boldsymbol{x}^{(k)}) - \nabla h(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k+0.5)}\rangle \\
&= \frac{1}{\eta}\left(\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) + \mathrm{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k+0.5)}) - \mathrm{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k)})\right) \enspace , && (4.6)
\end{aligned}$$

where the last line follows from Lemma 4.3.1.

On the other hand, by the first-order optimality condition at $\boldsymbol{x}^{(k+1)}$, for all $\boldsymbol{x} \in \mathcal{X}$,

$$\langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(k+1)} \rangle + \frac{1}{\eta} \langle \nabla h(\boldsymbol{x}^{(k+1)}) - \nabla h(\boldsymbol{x}^{(k)}), \boldsymbol{x} - \boldsymbol{x}^{(k+1)} \rangle \geq 0 \ .$$

Hence, for all $\boldsymbol{x} \in \mathcal{X}$,

$$\begin{aligned}\langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(k+1)} \rangle &\geq \frac{1}{\eta} \langle \nabla h(\boldsymbol{x}^{(k)}) - \nabla h(\boldsymbol{x}^{(k+1)}), \boldsymbol{x} - \boldsymbol{x}^{(k+1)} \rangle \\ &= \frac{1}{\eta} \left(\operatorname{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k)}) + \operatorname{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k+1)}) - \operatorname{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k)}) \right) \ ,\end{aligned}$$

where the last line once again follows from Lemma 4.3.1.

Continuing with the above inequality, for any given $\boldsymbol{x} \in \mathcal{X}$, we have:

$$\begin{aligned}&\frac{1}{\eta} \left(\operatorname{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k)}) + \operatorname{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k+1)}) - \operatorname{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k)}) \right) \\ &\leq \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(k+1)} \rangle \\ &= \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(k+0.5)} \rangle + \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k+1)} \rangle \\ &= \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(k+0.5)} \rangle + \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}) - \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k+1)} \rangle + \langle \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k+1)} \rangle \\ &\leq \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(k+0.5)} \rangle + \|\boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}) - \boldsymbol{f}(\boldsymbol{x}^{(k)})\| \cdot \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k+1)}\| + \langle \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k+1)} \rangle \ ,\end{aligned}$$

where the final line follows by the Cauchy-Schwarz inequality (Cauchy, 1821; Schwarz, 1884).

Recall that by the arithmetic mean-geometric mean inequality, $\forall x, y \in \mathbb{R}_+$, $\frac{x+y}{2} \geq \sqrt{xy}$. Hence, applying this inequality with $x = \eta \|\boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}) - \boldsymbol{f}(\boldsymbol{x}^{(k)})\|^2$ and $y = {}^1\!/\!_\eta \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k+1)}\|^2$

$$\begin{aligned}&\frac{1}{\eta} \left(\operatorname{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k)}) + \operatorname{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k+1)}) - \operatorname{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k)}) \right) \\ &\leq \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(k+0.5)} \rangle + \frac{\eta}{2} \|\boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}) - \boldsymbol{f}(\boldsymbol{x}^{(k)})\|^2 + \frac{1}{2\eta} \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k+1)}\|^2 + \langle \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k+1)} \rangle \\ &\leq \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(k+0.5)} \rangle + \eta \lambda^2 \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) + \frac{1}{2\eta} \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k+1)}\|^2 + \langle \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k+1)} \rangle \ ,\end{aligned}$$

where the last line was obtained by the assumption that there exists $\lambda \geq 0$, s.t. ${}^1\!/\!_2 \|\boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}) - \boldsymbol{f}(\boldsymbol{x}^{(k)})\|^2 \leq \lambda^2 \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})$.

Additionally, note that by the strong convexity of h, we have for all $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{X}$, $\operatorname{div}_h(\boldsymbol{x}, \boldsymbol{y}) \geq {}^1\!/\!_2 \|\boldsymbol{x} - \boldsymbol{y}\|^2$. Hence,

$$\begin{aligned}&\frac{1}{\eta} \left(\operatorname{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k)}) + \operatorname{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k+1)}) - \operatorname{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k)}) \right) \\ &\leq \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(k+0.5)} \rangle + \eta \lambda^2 \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) + \frac{1}{\eta} \operatorname{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k+0.5)}) + \langle \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k+1)} \rangle \ .\end{aligned}$$

Plugging Equation (4.6) into the above yields:

$$\frac{1}{\eta}\left(\mathrm{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k)}) + \mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k+1)}) - \mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k)})\right)$$
$$\leq \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(k+0.5)}\rangle + \eta\lambda^2\,\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})$$
$$+ \frac{1}{\eta}\mathrm{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k+0.5)}) - \frac{1}{\eta}\left(\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})\right.$$
$$\left.+\mathrm{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k+0.5)}) - \mathrm{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k)})\right)$$
$$\leq \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(k+0.5)}\rangle + \left(\eta\lambda^2 - \frac{1}{\eta}\right)\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) + \frac{1}{\eta}\mathrm{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k)}) \ .$$

Canceling out terms, the above inequality simplifies as:

$$\frac{1}{\eta}\left(\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k+1)}) - \mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k)})\right) \leq \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x} - \boldsymbol{x}^{(k+0.5)}\rangle +$$
$$\left(\eta\lambda^2 - \frac{1}{\eta}\right)\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) \ . \tag{4.7}$$

Finally, multiplying both sides by $-\eta < 0$, we obtain the statement of the lemma.

Lemma 4.3.2 already suggests convergence of the mirror extragradient method. To see this, suppose that the kernel function h is strictly convex, and that the algorithm has not yet converged, i.e., $\boldsymbol{x}^{(k+0.5)} \neq \boldsymbol{x}^{(k)}$. Under these assumptions, $\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) > 0$, and we can drop this term. Now, plugging in $\boldsymbol{x} = \boldsymbol{x}^*$ where $\boldsymbol{x}^* \in \mathcal{MVI}(\mathcal{X}, \boldsymbol{f})$, and rearranging yields:

$$\mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k)}) > \mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k+1)}) \ ,$$

implying asymptotic convergence of the mirror extragradient method.

While Lemma 4.3.2 implies asymptotic convergence to a strong solution, as the main theorem of this section of the thesis, we show polynomial-time computation of an ε-strong solution. To do so, it is necessary to bound the progress of the intermediate iterates $\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})$ as a function of the time horizon algorithm. In the proof of our main theorem, we show that we can bound this quantity, assuming the the VI satisfies the Minty condition, and that in addition, the kernel function is κ-Lipschitz-smooth.

Theorem 4.3.1 [Mirror Extragradient Method Convergence]**.**

Let $(\mathcal{X}, \boldsymbol{f})$ be a VI satisfying the Minty condition and let h be a 1-strongly-convex and κ-Lipschitz-smooth kernel function. Consider the mirror extragradient algorithm (Algorithm 3) run with the VI $(\mathcal{X}, \boldsymbol{f})$, the kernel function h, a step size $\eta > 0$, and a time horizon $\tau \in \mathbb{N}$. If there exists $\lambda \in (0, \frac{1}{\sqrt{2}\eta}]$, s.t. $\frac{1}{2}\left\|\boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}) - \boldsymbol{f}(\boldsymbol{x}^{(k)})\right\|^2 \leq \lambda^2 \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})$, then the sequence of outputs $\{\boldsymbol{x}^{(t+0.5)}, \boldsymbol{x}^{(t+1)}\}_t$ satisfies the following bound:

$$\min_{k=0,\dots,\tau} \max_{\boldsymbol{x} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x} \rangle \leq \frac{2(1+\kappa)\mathrm{diam}(\mathcal{X})}{\eta} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}} \quad ,$$

where $\boldsymbol{x}^* \in \mathcal{MVI}(\mathcal{X}, \boldsymbol{f})$ is a weak solution of the VI $(\mathcal{X}, \boldsymbol{f})$.

If, in addition, $\boldsymbol{x}^{(\tau)}_{\text{best}} \in \operatorname*{arg\,min}_{\boldsymbol{x}^{(k+0.5)}: k=0,\dots,\tau} \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})$, then for some choice of time horizon $\tau \in O(\frac{\kappa^2 \mathrm{diam}(\mathcal{X})^2 \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}{\eta^2 \varepsilon^2})$, $\boldsymbol{x}^{(\tau)}_{\text{best}}$ is an ε-strong solution of $(\mathcal{X}, \boldsymbol{f})$, and $\lim_{t\to\infty} \boldsymbol{x}^{(t+0.5)} = \lim_{t\to\infty} \boldsymbol{x}^{(t)} \in \mathcal{SVI}(\mathcal{X}, \boldsymbol{f})$ is a strong solution of the VI $(\mathcal{X}, \boldsymbol{f})$.

Proof of Theorem 4.3.1

Taking Lemma 4.3.2 with $\boldsymbol{x} \doteq \boldsymbol{x}^* \in \mathcal{MVI}(\mathcal{X}, \boldsymbol{f})$, we have:

$$\begin{aligned}
&\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k)}) - \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k+1)}) \\
&\geq \eta \underbrace{\langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^* \rangle}_{\geq 0} + \left(1 - (\eta\lambda)^2\right) \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) \\
&\geq \left(1 - (\eta\lambda)^2\right) \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})
\end{aligned}$$

Multiplying both sides by $\left(1 - (\eta\lambda)^2\right)^{-1} > 0$, we have:

$$\operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) \leq \frac{1}{1 - (\eta\lambda)^2} \left(\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k)}) - \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k+1)})\right) \tag{4.8}$$

Summing up over $k = 0, 1, \dots, \tau$, we have:

$$\begin{aligned}
\sum_{k=0}^{\tau} \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) &\leq \frac{1}{1 - (\eta\lambda)^2} \sum_{k=0}^{\tau} \left(\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k)}) - \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k+1)})\right) \\
&\leq \frac{1}{1 - (\eta\lambda)^2} \left(\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)}) - \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(\tau+1)})\right) \\
&\leq \frac{1}{1 - (\eta\lambda)^2} \left(\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})\right)
\end{aligned}$$

Dividing both sides by τ, we have:

$$\frac{1}{\tau}\sum_{k=0}^{\tau} \mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) \leq \frac{1}{\tau\,(1-(\eta\lambda)^2)}\left(\mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})\right)$$

$$\min_{k=0,\ldots,\tau} \mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) \leq \frac{1}{\tau\,(1-(\eta\lambda)^2)}\left(\mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})\right) \tag{4.9}$$

We can transform this convergence into a convergence in terms of a ε-strong solution. By the first-order optimality condition at $\boldsymbol{x}^{(k+0.5)}$, for all $\boldsymbol{x} \in \mathcal{X}$,

$$\langle \boldsymbol{f}(\boldsymbol{x}^{(k)}) + \frac{1}{\eta}\nabla h(\boldsymbol{x}^{(k+0.5)}) - \nabla h(\boldsymbol{x}^{(k)}), \boldsymbol{x} - \boldsymbol{x}^{(k+0.5)}\rangle \geq 0.$$

Re-organizing, for all $\boldsymbol{x} \in \mathcal{X}$ and $k \in \mathbb{N}$,

$$\begin{aligned}\langle \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}\rangle &\leq \frac{1}{\eta}\left\|\nabla h(\boldsymbol{x}^{(k+0.5)}) - \nabla h(\boldsymbol{x}^{(k)})\right\|\left\|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}\right\| \\ &\leq \frac{1}{\eta}\mathrm{diam}(\mathcal{X})\left\|\nabla h(\boldsymbol{x}^{(k+0.5)}) - \nabla h(\boldsymbol{x}^{(k)})\right\| \\ &\leq \frac{\kappa}{\eta}\mathrm{diam}(\mathcal{X})\left\|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)}\right\| \quad ,\end{aligned}$$

where the last line follow from the fact that h is κ-Lipschitz-smooth.

Now, with the above inequality in hand, for all $\boldsymbol{x} \in \mathcal{X}$ and $k \in \mathbb{N}$, we have:

$$\begin{aligned}&\langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}\rangle \\ &= \langle \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}\rangle + \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}) - \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}\rangle \\ &\leq \frac{\kappa}{\eta}\mathrm{diam}(\mathcal{X})\|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)}\| + \|\boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}) - \boldsymbol{f}(\boldsymbol{x}^{(k)})\| \cdot \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}\| \\ &\leq \frac{\kappa}{\eta}\mathrm{diam}(\mathcal{X})\|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)}\| + \lambda\sqrt{2\,\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})} \cdot \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}\| \\ &\leq \frac{\kappa}{\eta}\mathrm{diam}(\mathcal{X})\|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)}\| + \lambda\mathrm{diam}(\mathcal{X})\sqrt{2\,\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})} \\ &\leq \left(\frac{\kappa}{\eta} + \lambda\right)\mathrm{diam}(\mathcal{X})\sqrt{2\,\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})} \quad ,\end{aligned}$$

where the middle line was obtained by the assumption that there exists $\lambda \geq 0$, s.t. $^1/_2\|\boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}) - \boldsymbol{f}(\boldsymbol{x}^{(k)})\|^2 \leq \lambda^2\,\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})$, and the last line, from the strong convexity of h, which implies that $\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{y}) \geq {}^1/_2\|\boldsymbol{x} - \boldsymbol{y}\|^2$, or equivalently, $\sqrt{2\,\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{y})} \geq \|\boldsymbol{x} - \boldsymbol{y}\|^2$, for all $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{X}$.

Now, let $k^* \in \arg\min_{k=0,\dots,\tau} \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})$. We then have, for all $\boldsymbol{x} \in \mathcal{X}$,

$$\begin{aligned}\langle \boldsymbol{f}(\boldsymbol{x}^{(k^*+0.5)}), \boldsymbol{x}^{(k^*+0.5)} - \boldsymbol{x} \rangle &\leq \operatorname{diam}(\mathcal{X}) \left(\frac{\kappa}{\eta} + \lambda\right) \sqrt{2 \operatorname{div}_h(\boldsymbol{x}^{(k^*+0.5)}, \boldsymbol{x}^{(k^*)})} \\ &= \operatorname{diam}(\mathcal{X}) \left(\frac{\kappa}{\eta} + \lambda\right) \sqrt{2 \min_{k=0,\dots,\tau} \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})}\end{aligned}$$

Next, plugging Equation (4.9) into the above, we have, for all $\boldsymbol{x} \in \mathcal{X}$,

$$\langle \boldsymbol{f}(\boldsymbol{x}^{(k^*+0.5)}), \boldsymbol{x}^{(k^*+0.5)} - \boldsymbol{x} \rangle \leq \frac{\sqrt{2}\operatorname{diam}(\mathcal{X}) \left(\kappa/\eta + \lambda\right)}{\sqrt{1 - (\eta\lambda)^2}} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}}$$

Now, by the assumptions $\eta \leq \frac{1}{\sqrt{2}\lambda} \leq \frac{1}{\lambda}$, we have:

$$\begin{aligned}\langle \boldsymbol{f}(\boldsymbol{x}^{(k^*+0.5)}), \boldsymbol{x}^{(k^*+0.5)} - \boldsymbol{x} \rangle &\leq \frac{\sqrt{2}\operatorname{diam}(\mathcal{X}) \left(\kappa/\eta + 1/\eta\right)}{\sqrt{1 - (\eta\lambda)^2}} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}} \\ &= \frac{(1+\kappa)\sqrt{2}\operatorname{diam}(\mathcal{X})}{\eta\sqrt{1 - (\eta\lambda)^2}} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}} \\ &\leq \frac{(1+\kappa)\sqrt{2}\operatorname{diam}(\mathcal{X})}{\eta\sqrt{1 - (1/\sqrt{2})^2}} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}} \\ &\leq \frac{2(1+\kappa)\operatorname{diam}(\mathcal{X})}{\eta} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}}\end{aligned}$$

That is, we have:

$$\begin{aligned}\min_{k=0,\dots,\tau} \max_{\boldsymbol{x} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x} \rangle &\leq \max_{\boldsymbol{x} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}^{(k^*+0.5)}), \boldsymbol{x}^{(k^*+0.5)} - \boldsymbol{x} \rangle \\ &\leq \frac{2(1+\kappa)\operatorname{diam}(\mathcal{X})}{\eta} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}}\end{aligned}$$

In addition, for any $\varepsilon \geq 0$, letting $\frac{2(1+\kappa)\operatorname{diam}(\mathcal{X})}{\eta} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}} \leq \varepsilon$, and solving for τ, we have:

$$\begin{aligned}\frac{2(1+\kappa)\operatorname{diam}(\mathcal{X})}{\eta} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}} &\leq \varepsilon \\ \frac{4(1+\kappa)^2\operatorname{diam}(\mathcal{X})^2}{\eta^2} \frac{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}{\varepsilon^2} &\leq \tau\end{aligned}$$

That is, $\boldsymbol{x}^{(\tau)}_{\text{best}} \in \arg\min_{\boldsymbol{x}^{(k+0.5)}: k=0,\dots,\tau} \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})$ is an ε-strong solution after at most $\frac{4(1+\kappa)^2\operatorname{diam}(\mathcal{X})^2}{\eta^2} \frac{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}{\varepsilon^2}$ iterations of the mirror extragradient algorithm.

Finally, notice that

$$\lim_{k\to\infty} \max_{\boldsymbol{x} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x} \rangle = \lim_{\tau\to\infty} \min_{k=0,\dots,\tau} \max_{\boldsymbol{x} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x} \rangle = 0$$

and

$$\lim_{k\to\infty} \mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) = \lim_{\tau\to\infty} \min_{k=0,\dots,\tau} \mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) = 0$$

Hence, $\lim_{t\to\infty} \boldsymbol{x}^{(t+0.5)} = \lim_{t\to\infty} \boldsymbol{x}^{(t)} = \boldsymbol{x}^{**}$ is a strong solution of the VI $(\mathcal{X}, \boldsymbol{f})$.

With the main theorem of this section proven, some remarks are in order.

Remark 4.3.2.

First, note that the assumption that h is 1-strongly-convex is without loss of generality since any μ-strongly-convex kernel h' can be converted to a 1-strongly-convex kernel $1/\mu h'$.

Second, while for ease of exposition we assume that the VI is λ-Lipschitz-continuous, this assumption can more generally be weakened to 1) the VI is continuous and 2) for all $t \in \mathbb{N}$, there exists $\lambda \geq 0$ s.t. $\mathrm{div}_h(f(\boldsymbol{x}^{(t+0.5)}), f(\boldsymbol{x}^{(t)})) \leq \lambda \mathrm{div}_h(\boldsymbol{x}^{(t+0.5)}, \boldsymbol{x}^{(t)})$. The second part of this assumption can be interpreted as Lipschitz-continuity w.r.t. the Bregman divergence div_h over the trajectories of the mirror extragradient algorithm. As we will show in Chapter 6, this weaker assumption can be useful when the optimality operator does not satisfy Lipschitz-continuity.

4.3.3 Local Convergence of Mirror Extragradient

Unfortunately, beyond VIs for which the Minty condition holds, it seems implausible to devise a first-order method that converges to a strong solution. To see this, observe the following example.

Example 4.3.2 [Non-convergence in the absence of the Minty condition]**.**

Consider the VI $(\mathcal{X}, f)$ where $\mathcal{X} \doteq \mathbb{R}$ and $f(x) \doteq 1 - x^2$. A Minty solution fails to exist in this example (see, Example 4.1.2). In contrast, the set of strong solutions of $(\mathcal{X}, f)$ is given by $\mathcal{SVI}(\mathcal{X}, f) = \{-1, 1\}$. Notice that for any any $x > 1$, $f(x) < 0$. As such, for the mirror (extra)gradient method, for any choice of step size $\eta > 0$, if the initial iterate is initialized s.t. $x^{(0)} > 1$, then $x^{(k)} \to \infty$.

This is perhaps not surprising, since the computation of an ε-strong solution for Lipschitz-continuous VIs is in general a PPAD-complete problem (Kapron and Samieefar, 2024). Nevertheless, in the above example one can see that for $x^{(0)} < 1$, the mirror (extra)gradient algorithm converges to the strong solution $x^* = -1$. This then begs the question: *under what conditions can one guarantee the local convergence of the mirror extragradient algorithm to a strong solution?* As we will show, we can guarantee local convergence by assuming a local variant of Minty's condition. To define this condition, we rely on the notions of local weak and local strong solutions, recently introduced by Aussel and Chaipunya (2024).

Definition 4.3.3 [Local Weak and Strong Solution]**.**

Consider a VI $(\mathcal{X}, \mathcal{F})$ be a VI. Let $\delta \geq 0$ be a **locality parameter**.

A δ**-local strong solution** of the VI is an $\boldsymbol{x}^* \in \mathcal{X}$ that satisfies:

$$\exists \boldsymbol{f}(\boldsymbol{x}^*) \in \mathcal{F}(\boldsymbol{x}^*), \qquad \max_{\boldsymbol{x} \in \mathcal{X} \cap \mathcal{B}_\delta(\boldsymbol{x}^*)} \langle \boldsymbol{f}(\boldsymbol{x}^*), \boldsymbol{x} - \boldsymbol{x}^* \rangle \geq 0 \tag{4.10}$$

A δ**-local weak solution** of the VI is an $\boldsymbol{x}^* \in \mathcal{X}$ that satisfies:

$$\exists \boldsymbol{f}(\boldsymbol{x}) \in \mathcal{F}(\boldsymbol{x}), \qquad \max_{\boldsymbol{x} \in \mathcal{X} \cap \mathcal{B}_\delta(\boldsymbol{x}^*)} \langle \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x} - \boldsymbol{x}^* \rangle \leq 0 \tag{4.11}$$

We denote the set of δ-local strong (respectively, weak) solutions of a VI $(\mathcal{X}, \mathcal{F})$ by $\mathcal{LSVI}^\delta(\mathcal{X}, \mathcal{F})$ (respectively, $\mathcal{LMVI}^\delta(\mathcal{X}, \mathcal{F})$).

We note that for any VI $(\mathcal{X}, \mathcal{F})$ with $\mathcal{X}$ convex, local strong solutions are not of great interest since they coincide with (global) strong solutions (see section 3.2. of Aussel and Chaipunya (2024). This observation suggests that the computation of a δ-local strong solution is also PPAD-complete in Lipschitz-continuous VIs.

Nevertheless, as we will show next, local weak solutions can be of great interest to show local convergence to a strong solution. To understand how the above condition can imply convergence, recall by Lemma 4.3.2 the iterates of the mirror extragradient algorithm satisfy the following for all $t \in \mathbb{N}$:

$$\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k)}) - \mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k+1)}) \geq \eta \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x} \rangle + \left(1 - (\eta\lambda)^2\right) \mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})$$

Suppose that the kernel function h is strictly convex, and that the algorithm has not yet converged, i.e., $\boldsymbol{x}^{(k+0.5)} \neq \boldsymbol{x}^{(k)}$. Under these assumptions, $\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) > 0$, and we can drop this term. Re-organizing the expression, we then have:

$$\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k)}) > \mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k+1)}) + \eta\langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}\rangle$$

Now, notice that if we can ensure that for all $k \in \mathbb{N}_+$, there exists an $\boldsymbol{x}^* \in \mathcal{SVI}(\mathcal{X}, \boldsymbol{f})$ s.t. $\langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^*\rangle \geq 0$, then $\mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k)}) \geq \mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k+1)})$, implying that $\boldsymbol{x}^{(k)} \to \boldsymbol{x}^*$, i.e., $\boldsymbol{x}^{(k)}$ converges to a strong solution. If we were guaranteed the existence of a weak solution (i.e., the Minty condition), then we would plug in this weak solution for $\boldsymbol{x}$, which would guarantee that the second inner product term on the RHS would be equal to 0, in which case $\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k)}) > \mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(k+1)})$, and convergence to a weak solution (and thus assuming continuity of $\boldsymbol{f}$, also to a strong solution) would be guaranteed. As a weak solution might not exist, we will instead plug in a *local* weak solution for $\boldsymbol{x}$, and then choose our learning rate such that the mirror extragradient method remains close enough to this local weak solution throughout the run of the algorithm. Since any local weak solution is guaranteed to be a strong solution by Proposition 3.1 of Aussel and Chaipunya (2024), a weaker sufficient condition which ensures the existence of a strong solution $\boldsymbol{x}^* \in \mathcal{SVI}(\mathcal{X}, \boldsymbol{f})$ s.t. $\langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^*\rangle \geq 0$ is to initialize the algorithm with an initial iterate $\boldsymbol{x}^{(0)} \in \mathcal{X}$ that is $O(\delta)$-close to a δ-local weak solution $\boldsymbol{x}^* \in \mathcal{LMVI}^\delta(\mathcal{X}, \boldsymbol{f})$, for some $\delta \geq 0$, and in addition, to once again ensure that all subsequent intermediate iterates $\{\boldsymbol{x}^{(k+0.5)}\}_{k\in\mathbb{N}_{++}}$ remain δ-close to $\boldsymbol{x}^*$.

To ensure this property holds, we have to first bound the distance between the intermediate $\{\boldsymbol{x}^{(k+0.5)}\}_{k\in\mathbb{N}_+}$ and terminal $\{\boldsymbol{x}^{(k)}\}_{k\in\mathbb{N}_+}$ iterates. The following lemma provides us with such a bound.

Lemma 4.3.3 [Distance bound on intermediate iterates]**.**
Let $(\mathcal{X}, \boldsymbol{f})$ be a λ-Lipschitz-continuous VI satisfying the Minty condition, and h a 1-strongly-convex and κ-Lipschitz-smooth kernel function. Consider the mirror extragradient algorithm (Algorithm 3) run on the VI $(\mathcal{X}, \boldsymbol{f})$ with the kernel function h, any step size $\eta \geq 0$,

and for any time horizon $\tau \in \mathbb{N}$. The outputs $\{\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k+1)}\}_k$ satisfy for all $k \in [\tau]$,

$$\left\|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)}\right\| \leq \eta\ell \quad , \tag{4.12}$$

where $\ell \doteq \max_{\boldsymbol{x}\in\mathcal{X}}\|\boldsymbol{f}(\boldsymbol{x})\|$.

Proof of Lemma 4.3.3

Note that for all $k \in \mathbb{N}_+$, by the first-order optimality condition at $\boldsymbol{x}^{(k+0.5)}$, we have for all $\boldsymbol{x} \in \mathcal{X}$:

$$\langle \boldsymbol{f}(\boldsymbol{x}^{(k)}) + \frac{1}{\eta}\nabla h(\boldsymbol{x}^{(k+0.5)}) - \nabla h(\boldsymbol{x}^{(k)}), \boldsymbol{x} - \boldsymbol{x}^{(k+0.5)}\rangle \geq 0 \ .$$

Substituting $\boldsymbol{x} = \boldsymbol{x}^{(k)}$ above, we have:

$$\begin{aligned}\langle \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x}^{(k+0.5)}\rangle &\geq \frac{1}{\eta}\langle \nabla h(\boldsymbol{x}^{(k)}) - \nabla h(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k)} - \boldsymbol{x}^{(k+0.5)}\rangle \\ &= \frac{1}{\eta}\left(\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) + \mathrm{div}_h(\boldsymbol{x}^{(k)}, \boldsymbol{x}^{(k+0.5)})\right) \quad , \end{aligned} \tag{4.13}$$

where the last line follows from Lemma 4.3.1.

Re-organizing:

$$\begin{aligned}\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) &\leq \eta\langle \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x}^{(k+0.5)}\rangle - \mathrm{div}_h(\boldsymbol{x}^{(k)}, \boldsymbol{x}^{(k+0.5)}) & (4.14)\\ &\leq \eta\langle \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x}^{(k+0.5)}\rangle - \frac{1}{2}\left\|\boldsymbol{x}^{(k)} - \boldsymbol{x}^{(k+0.5)}\right\|^2 & (4.15)\\ &\leq \eta\left\|\boldsymbol{f}(\boldsymbol{x}^{(k)})\right\|\left\|\boldsymbol{x}^{(k)} - \boldsymbol{x}^{(k+0.5)}\right\| - \frac{1}{2}\left\|\boldsymbol{x}^{(k)} - \boldsymbol{x}^{(k+0.5)}\right\|^2 & (4.16)\\ &\leq \eta\ell\left\|\boldsymbol{x}^{(k)} - \boldsymbol{x}^{(k+0.5)}\right\| - \frac{1}{2}\left\|\boldsymbol{x}^{(k)} - \boldsymbol{x}^{(k+0.5)}\right\|^2 & (4.17)\end{aligned}$$

Since for all $z \in \mathbb{R}$ and $a, b \in \mathbb{R}_+$, it holds that $az - bz^2 \leq a^2/4b$, choosing $a = \eta\ell$ and $b = 1/2$ yields:

$$\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) \leq \frac{(\eta\ell)^2}{2} \tag{4.18}$$

Additionally, by the strong convexity of h, $\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{y}) \geq 1/2\|\boldsymbol{x} - \boldsymbol{y}\|^2$, or equivalently, $\sqrt{2\,\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{y})} \geq \|\boldsymbol{x} - \boldsymbol{y}\|^2$, for all $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{X}$. Hence,

$$\left\|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)}\right\| \leq \sqrt{2\,\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})} \leq \eta\ell \tag{4.19}$$

With the above Lemma in hand, we now show that if the initial iterate starts close enough to some local weak solution, then the intermediate iterates remain within this δ-ball for the duration of the run of the algorithm, for an appropriate choice of step size.

Lemma 4.3.4 [Mirror Extragradient Iterates Remain Local]**.**

Let $(\mathcal{X}, \boldsymbol{f})$ be a λ-Lipschitz-continuous VI satisfying the Minty condition and let h a 1-strongly-convex kernel function. Define $\ell \doteq \max_{\boldsymbol{x} \in \mathcal{X}} \|\boldsymbol{f}(\boldsymbol{x})\|$. Suppose that for some $\boldsymbol{x}^* \in \mathcal{LMVI}^\delta(\mathcal{X}, \boldsymbol{f})$ δ-local weak solution, the initial iterate $\boldsymbol{x}^{(0)} \in \mathcal{X}$ is chosen so that $\sqrt{2 \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})} \leq \delta - \eta\ell$. Consider the mirror extragradient algorithm (Algorithm 3) run on the VI $(\mathcal{X}, \boldsymbol{f})$ with the kernel function h, a step size $\eta \geq 0$, initial iterate $\boldsymbol{x}^{(0)}$, and some time horizon $\tau \in \mathbb{N}$. The outputs $\{\boldsymbol{x}^{(t+0.5)}, \boldsymbol{x}^{(t+1)}\}_t$ satisfy for all $t \in [\tau]$,

$$\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(t)}) \leq 1/2(\delta - \eta\ell)^2 \qquad \text{and} \left\|\boldsymbol{x}^{(t+0.5)} - \boldsymbol{x}^*\right\| \leq \delta \ .$$

Proof of Lemma 4.3.4

We prove the claim by induction on $t \in \mathbb{N}_+$.

Base case: $t = 0$

$$\begin{aligned}
\left\|\boldsymbol{x}^{(0.5)} - \boldsymbol{x}^*\right\| &= \left\|\boldsymbol{x}^{(0.5)} - \boldsymbol{x}^{(0)} + \boldsymbol{x}^{(0)} - \boldsymbol{x}^*\right\| \\
&\leq \left\|\boldsymbol{x}^{(0.5)} - \boldsymbol{x}^{(0)}\right\| + \left\|\boldsymbol{x}^{(0)} - \boldsymbol{x}^*\right\| \\
&\leq \left\|\boldsymbol{x}^{(0.5)} - \boldsymbol{x}^{(0)}\right\| + \sqrt{2 \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})} \\
&\leq \eta\ell + (\delta - \eta\ell) && \text{(Lemma 4.3.3)} \\
&\leq \delta
\end{aligned}$$

Inductive step: Suppose that for all $t = 0, \dots, \tau$, $\|\boldsymbol{x}^{(t+0.5)} - \boldsymbol{x}^*\| \leq \delta$ and $\sqrt{2 \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(t)})} \leq \delta - \eta\ell$ (or equivalently, $\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(t)}) \leq 1/2(\delta - \eta\ell)^2$). We will show that $\|\boldsymbol{x}^{(\tau+1.5)} - \boldsymbol{x}^*\| \leq \delta$ and $\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(\tau+1)}) \leq 1/2(\delta - \eta\ell)^2$.

By Lemma 4.3.2, we have:

$$\operatorname{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(\tau)}) - \operatorname{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(\tau+1)}) \tag{4.20}$$

$$\geq \eta\langle \boldsymbol{f}(\boldsymbol{x}^{(\tau+0.5)}), \boldsymbol{x}^{(\tau+0.5)} - \boldsymbol{x}\rangle + \left(1 - (\eta\lambda)^2\right) \operatorname{div}_h(\boldsymbol{x}^{(\tau+0.5)}, \boldsymbol{x}^{(\tau)}) \tag{4.21}$$

Substituting in $\boldsymbol{x} \doteq \boldsymbol{x}^* \in \mathcal{LMVI}^\delta(\mathcal{X}, \boldsymbol{f})$, we have:

$$\begin{aligned}
&\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(\tau)}) - \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(\tau+1)}) \\
&\geq \eta \underbrace{\left\langle \boldsymbol{f}(\boldsymbol{x}^{(\tau+0.5)}), \boldsymbol{x}^{(\tau+0.5)} - \boldsymbol{x}^* \right\rangle}_{\geq 0} + \left(1 - (\eta\lambda)^2\right) \underbrace{\operatorname{div}_h(\boldsymbol{x}^{(\tau+0.5)}, \boldsymbol{x}^{(\tau)})}_{\geq 0} \\
&\geq 0
\end{aligned}$$

Re-organizing, and using the induction hypothesis that $\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(\tau)}) \leq {}^1\!/_2(\delta - \eta\ell)^2$ yields:

$$ {}^1\!/_2(\delta - \eta\ell)^2 \geq \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(\tau)}) \geq \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(\tau+1)}) \tag{4.22}$$

Therefore,

$$\begin{aligned}
\left\|\boldsymbol{x}^{(\tau+0.5)} - \boldsymbol{x}^*\right\| &= \left\|\boldsymbol{x}^{(\tau+0.5)} - \boldsymbol{x}^{(\tau)} + \boldsymbol{x}^{(\tau)} - \boldsymbol{x}^*\right\| \\
&\leq \left\|\boldsymbol{x}^{(\tau+0.5)} - \boldsymbol{x}^{(\tau)}\right\| + \left\|\boldsymbol{x}^{(\tau)} - \boldsymbol{x}^*\right\| \\
&\leq \left\|\boldsymbol{x}^{(\tau+0.5)} - \boldsymbol{x}^{(\tau)}\right\| + \sqrt{2 \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(\tau)})} \\
&\leq \eta\ell + (\delta - \eta\ell) && \text{(Lemma 4.3.3)} \\
&\leq \delta
\end{aligned}$$

With the above lemma in hand, we can modify the proof of Theorem 4.3.1 slightly to show local convergence to a strong solution when the initial iterate of the algorithm is initialized close enough to a local solution.

Theorem 4.3.2 [Mirror Extragradient Method Local Convergence]**.**
Let $(\mathcal{X}, \boldsymbol{f})$ be a λ-Lipschitz-continuous VI, h a 1-strongly-convex and κ-Lipschitz-smooth kernel function, and let $\eta \in \left(0, \frac{1}{\sqrt{2}\lambda}\right]$. Suppose that for some $\boldsymbol{x}^* \in \mathcal{LMVI}^\delta(\mathcal{X}, \boldsymbol{f})$ δ-local weak solution, the initial iterate $\boldsymbol{x}^{(0)} \in \mathcal{X}$ is chosen so that $\sqrt{2 \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})} \leq \delta - \eta\ell$.

Consider the mirror extragradient algorithm (Algorithm 3) run on the VI $(\mathcal{X}, \boldsymbol{f})$ with the kernel function h, the step size η, initial iterate $\boldsymbol{x}^{(0)}$, and an arbitrary time horizon $\tau \in \mathbb{N}$.

The outputs $\{\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k+1)}\}_k$ satisfy for all $k \in [\tau]$,

$$\min_{k=0,\ldots,\tau} \max_{\boldsymbol{x}\in\mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}\rangle \leq \frac{\sqrt{2}(1+\kappa)\mathrm{diam}(\mathcal{X})}{\eta} \frac{\delta}{\sqrt{\tau}}$$

In addition, if $\boldsymbol{x}_{\text{best}}^{(\tau)} \in \arg\min_{\boldsymbol{x}^{(k+0.5)}:k=0,\ldots,\tau} \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)}\|$, then for some $\tau \in O(1/\varepsilon^2)$, $\boldsymbol{x}_{\text{best}}^{(\tau)}$ is an ε-strong solution of $(\mathcal{X}, \boldsymbol{f})$.

Proof of Theorem 4.3.2

Taking Lemma 4.3.2 with $\boldsymbol{x} \doteq \boldsymbol{x}^*$, where $\boldsymbol{x}^*$ is as given in the Theorem statement, by Lemma 4.3.4, for all $k \in \mathbb{N}$,

$$\begin{aligned}
&\mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k)}) - \mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k+1)}) \\
&\geq \eta \underbrace{\langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^*\rangle}_{\geq 0} + \left(1 - (\eta\lambda)^2\right) \mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) \\
&\geq \left(1 - (\eta\lambda)^2\right) \mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})
\end{aligned}$$

Multiplying both sides by $\left(1-(\eta\lambda)^2\right)^{-1} > 0$, we have:

$$\mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) \leq \frac{1}{1-(\eta\lambda)^2} \left(\mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k)}) - \mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k+1)})\right) \tag{4.23}$$

Summing up over $k = 0, 1, \ldots, \tau$, we have:

$$\sum_{k=0}^{\tau} \mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) \leq \frac{1}{1-(\eta\lambda)^2} \sum_{k=0}^{\tau} \left(\mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k)}) - \mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(k+1)})\right) \tag{4.24}$$

$$\leq \frac{1}{1-(\eta\lambda)^2} \left(\mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)}) - \mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(\tau+1)})\right) \tag{4.25}$$

$$\leq \frac{1}{1-(\eta\lambda)^2} \left(\mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})\right) \tag{4.26}$$

Dividing both sides by τ, we have:

$$\frac{1}{\tau}\sum_{k=0}^{\tau} \mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) \leq \frac{1}{\tau(1-(\eta\lambda)^2)} \left(\mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})\right) \tag{4.27}$$

$$\min_{k=0,\ldots,\tau} \mathrm{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)}) \leq \frac{1}{\tau(1-(\eta\lambda)^2)} \left(\mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})\right) \tag{4.28}$$

We can transform this convergence into a convergence in terms of a ε-strong solution. Now, recall by the first order optimality conditions of $\boldsymbol{x}^{(k+0.5)}$, we have for all $\boldsymbol{x} \in \mathcal{X}$:

$$\langle \boldsymbol{f}(\boldsymbol{x}^{(k)}) + \frac{1}{\eta}\langle \nabla h(\boldsymbol{x}^{(k+0.5)}) - \nabla h(\boldsymbol{x}^{(k)}), \boldsymbol{x} - \boldsymbol{x}^{(k+0.5)}\rangle \geq 0.$$

Re-organizing, for all $\boldsymbol{x} \in \mathcal{X}$, and $k \in \mathbb{N}$ we have:

$$\langle \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x} \rangle \leq \frac{1}{\eta} \left\| \nabla h(\boldsymbol{x}^{(k+0.5)}) - \nabla h(\boldsymbol{x}^{(k)}) \right\| \left\| \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x} \right\| \tag{4.29}$$

$$\leq \frac{\operatorname{diam}(\mathcal{X})}{\eta} \left\| \nabla h(\boldsymbol{x}^{(k+0.5)}) - \nabla h(\boldsymbol{x}^{(k)}) \right\| \tag{4.30}$$

$$\leq \frac{\operatorname{diam}(\mathcal{X})\kappa}{\eta} \left\| \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)} \right\| \tag{4.31}$$

where the last line follow from h being κ-Lipschitz-smooth.

Now, with the above inequality in hand, notice that for all $\boldsymbol{x} \in \mathcal{X}$ and $k \in \mathbb{N}$, we have:

$$\begin{aligned}
&\langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x} \rangle \\
&= \langle \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x} \rangle + \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}) - \boldsymbol{f}(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x} \rangle \\
&\leq \frac{\operatorname{diam}(\mathcal{X})\kappa}{\eta} \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)}\| + \|\boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}) - \boldsymbol{f}(\boldsymbol{x}^{(k)})\| \cdot \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}\| \\
&\leq \frac{\operatorname{diam}(\mathcal{X})\kappa}{\eta} \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)}\| + \lambda \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)}\| \cdot \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}\| \\
&\leq \operatorname{diam}(\mathcal{X}) \left(\frac{\kappa}{\eta} + \lambda \right) \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)}\|
\end{aligned}$$

where the penultimate line follows from the λ-Lipschitz-continuity of $\boldsymbol{f}$, and the strong convexity of h, which means that we have $\forall \boldsymbol{x}, \boldsymbol{y} \in \mathcal{X}$, $\operatorname{div}_h(\boldsymbol{x}, \boldsymbol{y}) \geq 1/2 \|\boldsymbol{x} - \boldsymbol{y}\|^2$..

Now, let $k^* \in \arg\min_{k=0,\ldots,\tau} \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)}\|$, we then have:

$$\begin{aligned}
\langle \boldsymbol{f}(\boldsymbol{x}^{(k^*+0.5)}), \boldsymbol{x}^{(k^*+0.5)} - \boldsymbol{x} \rangle &\leq \operatorname{diam}(\mathcal{X}) \left(\frac{\kappa}{\eta} + \lambda \right) \|\boldsymbol{x}^{(k^*+0.5)} - \boldsymbol{x}^{(k^*)}\| \\
&= \operatorname{diam}(\mathcal{X}) \left(\frac{\kappa}{\eta} + \lambda \right) \min_{k=0,\ldots,\tau} \|\boldsymbol{x}^{(k^*+0.5)} - \boldsymbol{x}^{(k^*)}\| \\
&\leq \operatorname{diam}(\mathcal{X}) \left(\frac{\kappa}{\eta} + \lambda \right) \min_{k=0,\ldots,\tau} \sqrt{2 \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})}
\end{aligned}$$

where the last line follows from

Now, plugging Equation (4.28) in the above, we have:

$$
\begin{aligned}
\langle \boldsymbol{f}(\boldsymbol{x}^{(k^*+0.5)}), \boldsymbol{x}^{(k^*+0.5)} - \boldsymbol{x}\rangle &\leq \operatorname{diam}(\mathcal{X}) \left(\frac{\kappa}{\eta} + \lambda\right) \min_{k=0,\dots,\tau} \sqrt{2 \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})} \\
&\leq \sqrt{2}\operatorname{diam}(\mathcal{X}) \left(\frac{\kappa}{\eta} + \lambda\right) \sqrt{\min_{k=0,\dots,\tau} \operatorname{div}_h(\boldsymbol{x}^{(k+0.5)}, \boldsymbol{x}^{(k)})} \\
&\leq \frac{\sqrt{2}\operatorname{diam}(\mathcal{X}) \left(\frac{\kappa}{\eta} + \lambda\right)}{\sqrt{1-(\eta\lambda)^2}} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}}
\end{aligned}
$$

Now, by the assumption that $\eta \leq \frac{1}{\sqrt{2}\lambda} \leq \frac{1}{\lambda}$, we have:

$$
\begin{aligned}
\langle \boldsymbol{f}(\boldsymbol{x}^{(k^*+0.5)}), \boldsymbol{x}^{(k^*+0.5)} - \boldsymbol{x}\rangle &\leq \frac{\sqrt{2}\operatorname{diam}(\mathcal{X}) \left(\frac{\kappa}{\eta} + \frac{1}{\eta}\right)}{\sqrt{1-(\eta\lambda)^2}} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}} \\
&= \frac{(1+\kappa)\sqrt{2}\operatorname{diam}(\mathcal{X})}{\eta\sqrt{1-(\eta\lambda)^2}} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}} \\
&= \frac{(1+\kappa)\sqrt{2}\operatorname{diam}(\mathcal{X})}{\eta\sqrt{1-(1/\sqrt{2})^2}} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}} \\
&= \frac{2(1+\kappa)\operatorname{diam}(\mathcal{X})}{\eta} \frac{\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})}}{\sqrt{\tau}} \\
&= \frac{\sqrt{2}(1+\kappa)\operatorname{diam}(\mathcal{X})}{\eta} \frac{\delta}{\sqrt{\tau}}
\end{aligned}
$$

where the last line follows from the assumption that $\sqrt{2 \operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})} \leq \delta - \eta\ell$ which implies $\sqrt{\operatorname{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})} \leq \frac{\delta}{\sqrt{2}}$.

That is, we have:

$$
\max_{\boldsymbol{x}\in\mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}^{(k+0.5)}), \boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}\rangle \leq \langle \boldsymbol{f}(\boldsymbol{x}^{(k^*+0.5)}), \boldsymbol{x}^{(k^*+0.5)} - \boldsymbol{x}\rangle \leq \frac{\sqrt{2}(1+\kappa)\operatorname{diam}(\mathcal{X})}{\eta} \frac{\delta}{\sqrt{\tau}}
$$

In addition, for any $\varepsilon \geq 0$, letting $\frac{\sqrt{2}(1+\kappa)\operatorname{diam}(\mathcal{X})}{\eta} \frac{\delta}{\sqrt{\tau}} \leq \varepsilon$, and solving for τ, we have:

$$
\begin{aligned}
\frac{\sqrt{2}(1+\kappa)\operatorname{diam}(\mathcal{X})}{\eta} \frac{\delta}{\sqrt{\tau}} &\leq \varepsilon \\
\frac{2(1+\kappa)^2\operatorname{diam}(\mathcal{X})^2}{\eta^2} \frac{\delta^2}{\varepsilon^2} &\leq \tau
\end{aligned}
$$

That is, $\boldsymbol{x}^{(\tau)}_{\text{best}} \in \arg\min_{\boldsymbol{x}^{(k+0.5)}:k=0,\dots,\tau} \|\boldsymbol{x}^{(k+0.5)} - \boldsymbol{x}^{(k)}\|$ is an ε-strong solution after at most $\frac{2(1+\kappa)^2\operatorname{diam}(\mathcal{X})^2}{\eta^2} \frac{\delta}{\varepsilon^2}$ iterations of the mirror extragradient algorithm.

4.4 Merit Function Methods

4.4.1 Merit Function Minimization via Second-Order Methods

Unfortunately, beyond domains where a local Minty solution might not exist, it is not possible to show even local convergence of the extragradient method for VIs. To see this clearly, consider the following example:

Example 4.4.1 [Non-convergence in the absence of local Minty solution]**.**
Consider the VI $(\mathcal{X}, \boldsymbol{f})$ where $\mathcal{X} \doteq \mathbb{R} \times \mathbb{R}$ and $\boldsymbol{f}(x, y) \doteq (x - y, x - y)$. The set of strong solutions of this VI is given by $\mathcal{SVI}(\mathcal{X}, \boldsymbol{f}) = \{(x, y) \in \mathcal{X} \mid x = y\}$. Notice that for any $x^{(0)} > y^{(0)}$, for any choice of step sizes, the iterates generated by the mirror (extra)gradient method tend to (positive) infinity, while for $x^{(0)} < y^{(0)}$, the iterates tend to negative infinity.

To overcome this non-convergence issue, given a VI $(\mathcal{X}, \boldsymbol{f})$, we will instead consider second-order methods. To derive a second-order method method, we will optimize a merit function associated with the VI.

Definition 4.4.1 [Merit functions]**.**
Given a VI $(\mathcal{X}, \boldsymbol{f})$. A function $\Xi : \mathcal{X} \to \mathbb{R}$ is said to be a **merit function** for the set of strong (respectively, weak) solutions of $(\mathcal{X}, \boldsymbol{f})$ iff

1. for all $\boldsymbol{x} \in \mathcal{X}$, $\Xi(\boldsymbol{x}) \geq 0$
2. $\arg\min_{\boldsymbol{x} \in \mathcal{X}} \Xi(\boldsymbol{x}) = \mathcal{SVI}(\mathcal{X}, \boldsymbol{f})$ (respectively, $\arg\min_{\boldsymbol{x} \in \mathcal{X}} \Xi(\boldsymbol{x}) = \mathcal{MVI}(\mathcal{X}, \boldsymbol{f})$)

The canonical merit function associated with the strong solution of any VI is the **primal gap function** $\Xi : \mathcal{X} \to \mathbb{R}_+$:

$$\Xi(\boldsymbol{x}) \doteq \max_{\boldsymbol{y} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x} - \boldsymbol{y} \rangle \tag{4.32}$$

Analogously, we can also define a merit function $\chi : \mathcal{X} \to \mathbb{R}_+$, known as the **dual gap function**, which is associated with weak solutions:

$$\chi(\boldsymbol{x}) \doteq \max_{\boldsymbol{y} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{y}), \boldsymbol{x} - \boldsymbol{y} \rangle \tag{4.33}$$

An important property of the primal (respectively, dual) gap function is that its set of minima coincide with the set of strong (respectively, weak) solutions, i.e., $\mathcal{SVI}(\mathcal{X}, \boldsymbol{f}) = \arg\min_{\boldsymbol{x} \in \mathcal{X}} \Xi(\boldsymbol{x})$ (respectively, $\mathcal{MVI}(\mathcal{X}, \boldsymbol{f}) = \arg\min_{\boldsymbol{x} \in \mathcal{X}} \chi(\boldsymbol{x})$) (see, for instance, Proposition 2.3 and 2.4 of Huang and Zhang (2023)). In other words, the primal (respectively, dual) gap function is a merit function for the strong (respectively, weak) solutions of the VI. While the primal gap function is in general non-convex and non-differentiable, the dual gap function is always convex. Its evaluation, however, requires solving a non-convex optimization problem. As such, finding even stationary points of the primal and dual gap functions is in general intractable.

Nevertheless, it is possible to formulate a differentiable merit function for strong solutions called the α**-regularized primal function** $\Xi_\alpha : \mathcal{X} \to \mathbb{R}_+$:

$$\Xi_\alpha(\boldsymbol{x}) \doteq \max_{\boldsymbol{y} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x} - \boldsymbol{y} \rangle - \frac{\alpha}{2} \|\boldsymbol{y} - \boldsymbol{x}\|^2 \quad , \tag{4.34}$$

where $\alpha > 0$ is a regularization parameter.

We note the following important lemma due to Fukushima (1992), which we prove here for completeness.

Lemma 4.4.1 [Properties of the regularized primal gap]**.**
Consider a continuous VI $(\mathcal{X}, \boldsymbol{f})$. If $\alpha > 0$, then $\max_{\boldsymbol{y} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x} - \boldsymbol{y} \rangle - \alpha/2 \|\boldsymbol{y} - \boldsymbol{x}\|^2$ has a unique solution. In addition, the following hold:

1. $\boldsymbol{y}^*(\boldsymbol{x}) = \arg\max_{\boldsymbol{y} \in \mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x} - \boldsymbol{y} \rangle - \frac{\alpha}{2} \|\boldsymbol{y} - \boldsymbol{x}\|^2 \doteq \Pi_{\mathcal{X}} \left[\boldsymbol{x} - \frac{1}{\alpha} \boldsymbol{f}(\boldsymbol{x})\right]$

2. $\nabla \Xi_\alpha(\boldsymbol{x}) = \boldsymbol{f}(\boldsymbol{x}) - (\nabla \boldsymbol{f}(\boldsymbol{x}) + \alpha \mathbb{I}) (\boldsymbol{y}^*(\boldsymbol{x}) - \boldsymbol{x})$

3. $\Xi_\alpha(\boldsymbol{x}) = \max_{\boldsymbol{y} \in \mathcal{X}} \frac{\alpha}{2} \left(\left\| \frac{1}{\alpha} \boldsymbol{f}(\boldsymbol{x}) \right\|^2 - \left\| \boldsymbol{y} - \left(\boldsymbol{x} - \frac{1}{\alpha} \boldsymbol{f}(\boldsymbol{x})\right) \right\|^2 \right)$

4. For all $\boldsymbol{x} \in \mathcal{X}$, $\Xi_\alpha(\boldsymbol{x}) \geq 0$ and $\mathcal{SVI}(\mathcal{X}, \boldsymbol{f}) = \arg\min_{\boldsymbol{x} \in \mathcal{X}} \Xi_\alpha(\boldsymbol{x})$

Proof of Lemma 4.4.1

For the first part, by the first-order optimality condition, for all $\boldsymbol{x} \in \mathcal{X}$,

$$-\boldsymbol{f}(\boldsymbol{x}) - \alpha\left(\boldsymbol{y}^* - \boldsymbol{x}\right) \in \left\{\boldsymbol{z}' \in \mathcal{X} \mid \left\langle \boldsymbol{z}', \boldsymbol{z} - \boldsymbol{y}^* \right\rangle \geq 0, \forall \boldsymbol{z} \in \mathcal{X}\right\}$$

Re-organizing, we have for all $\boldsymbol{x} \in \mathcal{X}$:

$$\begin{aligned}
\boldsymbol{y}^* &\in \boldsymbol{x} - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}) + \left\{\boldsymbol{z}' \mid \left\langle \boldsymbol{z}', \boldsymbol{z} - \boldsymbol{y}^* \right\rangle \geq 0 \mid \forall \boldsymbol{z}', \boldsymbol{z} \in \mathcal{X}\right\} \\
&\in \left\{\boldsymbol{x} - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}) + \boldsymbol{z}' \mid \left\langle \boldsymbol{z}', \boldsymbol{z} - \boldsymbol{y}^* \right\rangle \geq 0 \mid \forall \boldsymbol{z}', \boldsymbol{z} \in \mathcal{X}\right\} \\
&\in \left\{\boldsymbol{y}^* \in \mathcal{X} \mid \left\langle \boldsymbol{y}^* - (\boldsymbol{x} - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})), \boldsymbol{z} - \boldsymbol{y}^* \right\rangle \geq 0 \mid \forall \boldsymbol{z} \in \mathcal{X}\right\} \\
&\in \left\{\boldsymbol{y}^* \in \mathcal{X} \mid \left\langle (\boldsymbol{x} - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})) - \boldsymbol{y}^*, \boldsymbol{z} - \boldsymbol{y}^* \right\rangle \leq 0 \mid \forall \boldsymbol{z} \in \mathcal{X}\right\} \\
&\in \operatorname*{arg\,min}_{\boldsymbol{y} \in \mathcal{X}} \left\| \boldsymbol{x} - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{y} \right\|^2
\end{aligned}$$

For the second part of the lemma, applying Danskin's Theorem (Danskin, 1966), we have:

$$\begin{aligned}
\nabla \Xi_\alpha(\boldsymbol{x}) &\doteq \boldsymbol{f}(\boldsymbol{x}) + \langle \nabla \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x} - \boldsymbol{y}^*(\boldsymbol{x}) \rangle - \alpha\left(\boldsymbol{y}^*(\boldsymbol{x}) - \boldsymbol{x}\right) \\
&= \boldsymbol{f}(\boldsymbol{x}) - \langle \nabla \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{y}^*(\boldsymbol{x}) - \boldsymbol{x} \rangle - \alpha\left(\boldsymbol{y}^*(\boldsymbol{x}) - \boldsymbol{x}\right) \\
&= \boldsymbol{f}(\boldsymbol{x}) - \left(\nabla \boldsymbol{f}(\boldsymbol{x}) + \alpha \mathbb{I}\right)\left(\boldsymbol{y}^*(\boldsymbol{x}) - \boldsymbol{x}\right)
\end{aligned}$$

For the third part of the lemma, we have:

$$
\begin{aligned}
&\Xi_\alpha(\boldsymbol{x})\\
&= \max_{\boldsymbol{y}\in\mathcal{X}} \langle \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x}-\boldsymbol{y}\rangle - \frac{\alpha}{2}\|\boldsymbol{y}-\boldsymbol{x}\|^2\\
&= \langle \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})\rangle - \frac{\alpha}{2}\|\boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})\|^2\\
&= \langle \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})\rangle - \frac{\alpha}{2}\langle \boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x}), \boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})\rangle\\
&= \left\langle \boldsymbol{f}(\boldsymbol{x}) - \frac{\alpha}{2}(\boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})), \boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})\right\rangle\\
&= \frac{\alpha}{2}\left\langle \frac{2}{\alpha}\boldsymbol{f}(\boldsymbol{x}) - (\boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})), \boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})\right\rangle\\
&= \frac{\alpha}{2}\left[\left\langle \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})\right\rangle + \left\langle \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}) - (\boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})), \boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})\right\rangle\right]\\
&= \frac{\alpha}{2}\left[\left\langle \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})\right\rangle - \left\langle \left(\boldsymbol{x}-\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right) - \boldsymbol{y}^*(\boldsymbol{x}), \boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})\right\rangle\right]\\
&= \frac{\alpha}{2}\left[\left\langle \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})\right\rangle - \left\langle \left(\boldsymbol{x}-\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right) - \boldsymbol{y}^*(\boldsymbol{x}), \boldsymbol{x}-\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})+\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})-\boldsymbol{y}^*(\boldsymbol{x})\right\rangle\right]\\
&= \frac{\alpha}{2}\left[\left\langle \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x}-\boldsymbol{y}^*(\boldsymbol{x})\right\rangle - \left\|\left(\boldsymbol{x}-\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right) - \boldsymbol{y}^*(\boldsymbol{x})\right\|^2 - \left\langle \left(\boldsymbol{x}-\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right) - \boldsymbol{y}^*(\boldsymbol{x}), \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right\rangle\right]\\
&= \frac{\alpha}{2}\left[\left\langle \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}), \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right\rangle - \left\|\left(\boldsymbol{x}-\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right) - \boldsymbol{y}^*(\boldsymbol{x})\right\|^2\right]\\
&= \frac{\alpha}{2}\left[\left\|\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right\|^2 - \left\|\left(\boldsymbol{x}-\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right) - \boldsymbol{y}^*(\boldsymbol{x})\right\|^2\right]\\
&= \max_{\boldsymbol{y}\in\mathcal{X}} \frac{\alpha}{2}\left[\left\|\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right\|^2 - \left\|\boldsymbol{y} - \left(\boldsymbol{x}-\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right)\right\|^2\right]\\
&= \frac{\alpha}{2}\left[\left\|\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right\|^2 - \left\|\boldsymbol{y}^*(\boldsymbol{x}) - \left(\boldsymbol{x}-\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right)\right\|^2\right]
\end{aligned}
$$

For the final part, first note that we have by the third part of the lemma, for all $\boldsymbol{x}\in\mathcal{X}$,

$$
\begin{aligned}
\Xi_\alpha(\boldsymbol{x}) &= \max_{\boldsymbol{y}\in\mathcal{X}} \frac{\alpha}{2}\left[\left\|\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right\|^2 - \left\|\boldsymbol{y} - (\boldsymbol{x}-\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}))\right\|^2\right]\\
&\geq \frac{\alpha}{2}\left[\left\|\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right\|^2 - \left\|\boldsymbol{x} - \left(\boldsymbol{x}-\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right)\right\|^2\right] = 0 \ ,
\end{aligned}
$$

where the final equality follows from the fact that $\|{}^1\!/\!_\alpha \boldsymbol{f}(\boldsymbol{x})\|^2 = \|(\boldsymbol{x} - {}^1\!/\!_\alpha \boldsymbol{f}(\boldsymbol{x})) - \boldsymbol{x}\|^2$.

Hence, we can have $\Xi_\alpha(\boldsymbol{x}) = 0$ iff $\boldsymbol{y}^*(\boldsymbol{x}) = \boldsymbol{x}$, which, by the definition of $\boldsymbol{y}^*(\boldsymbol{x})$, implies $\Xi(\boldsymbol{x}) = \max_{\boldsymbol{y}\in\mathcal{X}}\langle \boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x}-\boldsymbol{y}\rangle = 0$. That is, $\mathcal{SVI}(\mathcal{X}, \boldsymbol{f}) = \arg\min_{\boldsymbol{x}\in\mathcal{X}} \Xi_\alpha(\boldsymbol{x})$, proving the final part of the lemma.

The first and second part of this lemma show that the gradient of the regularized primal gap function Ξ_α can be evaluated with a constant number of evaluations of the optimality operator $\boldsymbol{f}$ of the VI and its gradient $\nabla \boldsymbol{f}$. Hence, we can minimize the regularized primal gap function, at least locally, using a gradient descent method. Importantly, as this gradient depends on $\nabla \boldsymbol{f}$, the ensuing algorithm, which we call the **mirror potential algorithm** is a second-order method (Algorithm 5).

Algorithm 5 Mirror Potential Algorithm

Input: $\Xi_\alpha, h, \tau, \eta, \boldsymbol{x}^{(0)}$

Output: $\{\boldsymbol{x}^{(t)}\}_{t \in [\tau]}$

1: **for** $t = 1, \dots, \tau$ **do**

2: $\quad \boldsymbol{x}^{(t+1)} \leftarrow \arg\min_{\boldsymbol{x} \in \mathcal{X}} \left\{ \left\langle \nabla \Xi_\alpha(\boldsymbol{x}^{(t)}), \boldsymbol{x} - \boldsymbol{x}^{(t)} \right\rangle + \frac{1}{2\eta} \mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}^{(t)}) \right\}$

return $\{\boldsymbol{x}^{(t)}\}_{t \in [\tau]}$

Next, to prove an asymptotic convergence bound for the mirror potential method, we require that the regularized primal gap function Ξ_α is at a minimum a weakly-concave function. It turns out that when the optimality operator $\boldsymbol{f}$ is Lipschitz-continuous *and* Lipschitz-smooth, the regularized primal gap function is weakly-concave. To prove this, we first prove the following technical lemma, which is a slight variant of Lemma 4.2 of Drusvyatskiy and Paquette (2019) for function composition that yields weakly-convex functions.

Lemma 4.4.2 [Composition of weakly-concave and Lipschitz-smooth functions]**.**
Given an ℓ-Lipschitz-continuous and ρ-weakly-concave function $h : \mathcal{X} \to \mathbb{R}$, and a λ-Lipschitz-continuous and β-Lipschitz-smooth function $\boldsymbol{c} : \mathcal{X} \to \mathcal{X}$. Their composition $\boldsymbol{x} \mapsto h(\boldsymbol{c}(\boldsymbol{x}))$ is $\beta\ell + \rho\lambda^2$-weakly-concave.

Proof of Lemma 4.4.2

Let $\varphi(\boldsymbol{x}) \doteq h(\boldsymbol{c}(\boldsymbol{x}))$.

First, note the following:

$$\begin{aligned}
&\|\nabla h(\boldsymbol{c}(\boldsymbol{x}))[\boldsymbol{c}(\boldsymbol{y}) - \boldsymbol{c}(\boldsymbol{x}))] - \nabla h(\boldsymbol{c}(\boldsymbol{x}))[\nabla \boldsymbol{c}(\boldsymbol{x})(\boldsymbol{y} - \boldsymbol{x})]\|^2 \\
&= \|\nabla h(\boldsymbol{c}(\boldsymbol{x}))[\boldsymbol{c}(\boldsymbol{y}) - \boldsymbol{c}(\boldsymbol{x})) - \nabla \boldsymbol{c}(\boldsymbol{x})(\boldsymbol{y} - \boldsymbol{x})]\|^2 \\
&\leq \|\nabla h(\boldsymbol{c}(\boldsymbol{x}))\|^2 \|\boldsymbol{c}(\boldsymbol{y}) - \boldsymbol{c}(\boldsymbol{x}) - \nabla \boldsymbol{c}(\boldsymbol{x})(\boldsymbol{y} - \boldsymbol{x})\|^2 \\
&\leq \ell^2 \|\boldsymbol{c}(\boldsymbol{y}) - \boldsymbol{c}(\boldsymbol{x}) - \nabla \boldsymbol{c}(\boldsymbol{x})(\boldsymbol{y} - \boldsymbol{x})\|^2 \\
&\leq \frac{\beta \ell^2}{2} \|\boldsymbol{y} - \boldsymbol{x}\|^2 \ ,
\end{aligned}$$

where the penultimate line follows from the Lipschitz-continuity of h, and the final line, by the β-Lipschitz-smoothness of $\boldsymbol{c}$, which ensures that $\|\boldsymbol{c}(\boldsymbol{y}) - \boldsymbol{c}(\boldsymbol{x}) - \nabla \boldsymbol{c}(\boldsymbol{x})(\boldsymbol{y} - \boldsymbol{x})\| \leq \beta/2 \|\boldsymbol{y} - \boldsymbol{x}\|^2$.

Combining the above with the ρ-weak-concavity of h, which ensures that $h(\boldsymbol{z}) \leq h(\boldsymbol{z}') + \langle \nabla h(\boldsymbol{z}'), \boldsymbol{z} - \boldsymbol{z}' \rangle + \rho/2 \, \|\boldsymbol{z} - \boldsymbol{z}'\|$, yields:

$$\begin{aligned}
\varphi(\boldsymbol{y}) = h(\boldsymbol{c}(\boldsymbol{y})) &\leq h(\boldsymbol{c}(\boldsymbol{x})) + \langle \nabla h(\boldsymbol{c}(\boldsymbol{x})), \boldsymbol{c}(\boldsymbol{y}) - \boldsymbol{c}(\boldsymbol{x}) \rangle + \frac{\rho}{2} \|\boldsymbol{c}(\boldsymbol{y}) - \boldsymbol{c}(\boldsymbol{x})\|^2 \\
&\leq \varphi(\boldsymbol{x}) + \langle \nabla h(\boldsymbol{c}(\boldsymbol{x})), \boldsymbol{c}(\boldsymbol{y}) - \boldsymbol{c}(\boldsymbol{x}) \rangle + \frac{\rho \lambda^2}{2} \|\boldsymbol{y} - \boldsymbol{x}\|^2 \\
&\leq \varphi(\boldsymbol{x}) + \langle \nabla h(\boldsymbol{c}(\boldsymbol{x})), \nabla \boldsymbol{c}(\boldsymbol{x})(\boldsymbol{y} - \boldsymbol{x}) \rangle^2 + \frac{\beta \ell}{2} \|\boldsymbol{y} - \boldsymbol{x}\|^2 + \frac{\rho \lambda^2}{2} \|\boldsymbol{y} - \boldsymbol{x}\|^2 \\
&= \varphi(\boldsymbol{x}) + \langle \nabla \boldsymbol{c}(\boldsymbol{x}) \nabla h(\boldsymbol{c}(\boldsymbol{x})), \boldsymbol{y} - \boldsymbol{x} \rangle^2 + \frac{\beta \ell}{2} \|\boldsymbol{y} - \boldsymbol{x}\|^2 + \frac{\rho \lambda^2}{2} \|\boldsymbol{y} - \boldsymbol{x}\|^2 \\
&= \varphi(\boldsymbol{x}) + \langle \nabla \varphi(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle^2 + \frac{\beta \ell + \rho \lambda^2}{2} \|\boldsymbol{y} - \boldsymbol{x}\|^2
\end{aligned}$$

With the above lemma in hand, we can now define a class of VIs for which the regularized primal gap function is weakly-concave. To this end, we define the class of Lipschitz-smooth of VIs.

Definition 4.4.2 [Lipschitz-Smooth VIs]**.**
Given a modulus of smoothness $\beta \geq 0$, a VI $(\mathcal{X}, \boldsymbol{f})$ is β-**Lipschitz-smooth** iff $\boldsymbol{f}$ is β-Lipschitz-smooth.

Next, using Lemma 4.4.2 and the above definition, we prove that the regularized primal gap function is weakly-concave in Lipschitz-continuous and Lipschitz-smooth VIs.

Lemma 4.4.3 [Weak-concavity of regularized primal gap]**.**

Consider a λ-Lipschitz-continuous and β-Lipschitz-smooth VI $(\mathcal{X}, \boldsymbol{f})$. Then, for all $\alpha \geq 0$, the regularized primal gap function Ξ_α associated with $(\mathcal{X}, \boldsymbol{f})$ is $(2\beta\alpha \mathrm{diam}(\mathcal{X})^2 + 1 + 2\lambda)$-weakly-concave.

Proof of Lemma 4.4.3

Let $h(\boldsymbol{z}, \boldsymbol{z}'; \boldsymbol{y}) \doteq \frac{\alpha}{2}\left[\|\boldsymbol{z}\|^2 - \|\boldsymbol{y} - \boldsymbol{z}'\|^2\right]$. Notice that for all $\boldsymbol{y} \in \mathcal{X}$, $h(\boldsymbol{z}, \boldsymbol{z}'; \boldsymbol{y}) + \frac{\alpha}{2}\|\boldsymbol{z}\|^2 = -\|\boldsymbol{y} - \boldsymbol{z}'\|^2$, and hence $(\boldsymbol{z}, \boldsymbol{z}') \mapsto h(\boldsymbol{z}, \boldsymbol{z}'; \boldsymbol{y}) + \frac{\alpha}{2}\|\boldsymbol{z}\|^2$ is concave. That is, for all $\boldsymbol{y} \in \mathcal{X}$, $h(\boldsymbol{z}, \boldsymbol{z}'; \boldsymbol{y})$ is α-weakly-concave. In addition, h is $\alpha\mathrm{diam}(\mathcal{X})$-Lipschitz-continuous, since $\|\nabla_{\boldsymbol{z},\boldsymbol{z}'} h(\boldsymbol{z}, \boldsymbol{z}'; \boldsymbol{y})\| = \alpha\|(\boldsymbol{z}, \boldsymbol{y} - \boldsymbol{z}')\| \leq 2\alpha\mathrm{diam}(\mathcal{X})$

Let $\boldsymbol{c}(\boldsymbol{x}) \doteq (\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}), \boldsymbol{x} - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}))$. Now, notice that $\boldsymbol{c}$ is $\frac{1+2\lambda}{\alpha}$-Lipschitz-continuous since

$$\left\|\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}) - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{y})\right\| \leq \frac{1}{\alpha}\|\boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{f}(\boldsymbol{y})\| \leq \frac{\lambda}{\alpha}\|\boldsymbol{x} - \boldsymbol{y}\|$$

and

$$\left\|\boldsymbol{x} - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{y} - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{y})\right\| \tag{4.35}$$

$$= \left\|\boldsymbol{x} - \boldsymbol{y} + \frac{1}{\alpha}(\boldsymbol{f}(\boldsymbol{y}) - \boldsymbol{f}(\boldsymbol{x}))\right\| \tag{4.36}$$

$$\leq \|\boldsymbol{x} - \boldsymbol{y}\| + \frac{1}{\alpha}\|\boldsymbol{f}(\boldsymbol{y}) - \boldsymbol{f}(\boldsymbol{x})\| \tag{4.37}$$

$$\leq \|\boldsymbol{x} - \boldsymbol{y}\| + \frac{\lambda}{\alpha}\|\boldsymbol{y} - \boldsymbol{x}\| \tag{4.38}$$

$$\leq \frac{1+\lambda}{\alpha}\|\boldsymbol{x} - \boldsymbol{y}\| \tag{4.39}$$

In other words, $\boldsymbol{c}$ is $\frac{1+2\lambda}{\alpha}$-Lipschitz-continuous since $\|\boldsymbol{c}(\boldsymbol{x}) - \boldsymbol{c}(\boldsymbol{y})\| \leq \left\|\frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}) - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{y})\right\| + \left\|\boldsymbol{x} - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x}) - \boldsymbol{y} - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{y})\right\| \leq \frac{1+2\lambda}{\alpha}\|\boldsymbol{x} - \boldsymbol{y}\|$.

Additionally, notice that $\boldsymbol{x} \mapsto \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})$ is $\frac{\beta}{\alpha}$-Lipschitz-smooth, since $\left\|\frac{1}{\alpha}\nabla\boldsymbol{f}(\boldsymbol{x}) - \frac{1}{\alpha}\nabla\boldsymbol{f}(\boldsymbol{y})\right\| \leq \frac{1}{\alpha}\|\nabla\boldsymbol{f}(\boldsymbol{x}) - \nabla\boldsymbol{f}(\boldsymbol{y})\| \leq \frac{\beta}{\alpha}\|\boldsymbol{x} - \boldsymbol{y}\|$. Similarly, $\boldsymbol{x} \mapsto \boldsymbol{x} - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})$ is $\frac{\beta}{\alpha}$-Lipschitz-smooth, since since $\left\|1 - \frac{1}{\alpha}\nabla\boldsymbol{f}(\boldsymbol{x}) - 1 + \frac{1}{\alpha}\nabla\boldsymbol{f}(\boldsymbol{y})\right\| \leq$

$\frac{1}{\alpha}\|\nabla \boldsymbol{f}(\boldsymbol{x}) - \nabla \boldsymbol{f}(\boldsymbol{y})\| \leq \frac{\beta}{\alpha}\|\boldsymbol{x} - \boldsymbol{y}\|$. As a result, $\boldsymbol{c}$ is $\frac{2\beta}{\alpha}$-Lipschitz-smooth since $\|\nabla \boldsymbol{c}(\boldsymbol{x}) - \nabla \boldsymbol{c}(\boldsymbol{y})\| \leq \left\|\frac{1}{\alpha}\nabla \boldsymbol{f}(\boldsymbol{x}) - \frac{1}{\alpha}\nabla \boldsymbol{f}(\boldsymbol{y})\right\| + \left\|1 - \frac{1}{\alpha}\nabla \boldsymbol{f}(\boldsymbol{x}) - 1 + \frac{1}{\alpha}\nabla \boldsymbol{f}(\boldsymbol{y})\right\| \leq \frac{2\beta}{\alpha}\|\boldsymbol{x} - \boldsymbol{y}\|$.

Hence, by Lemma 4.4.2, $\boldsymbol{x} \mapsto h(\boldsymbol{c}(\boldsymbol{x}))$ is $\frac{2\beta}{\alpha}\alpha^2 \mathrm{diam}(\mathcal{X})^2 + \frac{1+2\lambda}{\alpha}\alpha = (2\beta\alpha\mathrm{diam}(\mathcal{X})^2 + 1 + 2\lambda)$-weakly-concave, since it is the composition of h, which is α-weakly-concave and $\alpha\mathrm{diam}(\mathcal{X})$-Lipschitz-continuous, and $\boldsymbol{c}$ which is $\frac{1+2\lambda}{\alpha}$-Lipschitz-continuous and $\frac{2\beta}{\alpha}$-Lipschitz-smooth.

4.4.2 Mirror Potential Algorithm for VIs

With Lemma 4.4.3 in hand, we can now analyze the convergence properties of the mirror potential algorithm. Since Ξ_α is in general non-convex, it is PPAD complete to compute a minimum of Ξ_α in Lipschitz-continuous and Lipschitz-smooth VIs, because otherwise we would have computed a strong solution of such a VI, which is a PPAD-complete problem (Kapron and Samieefar, 2024). We will thus aim to compute a stationary point of Ξ_α (Definition 4.4.3) instead.

Definition 4.4.3 [Stationary point]**.**
Given an optimization problem $\min_{\boldsymbol{x}\in\mathcal{X}} h(\boldsymbol{x})$ where $\mathcal{X} \subseteq \mathcal{U}$ is the constraint set and $h : \mathcal{U} \to \mathbb{R}$ is a subdifferentiable objective, for any approximation parameter $\varepsilon \geq 0$, an ε-**stationary point** $\boldsymbol{x}^* \in \mathcal{X}$ is defined as the set of ε-strong solutions $\mathcal{SVI}(\mathcal{X}, \mathcal{D}h)$ of the VI $(\mathcal{X}, \mathcal{D}h)$.

A 0-stationary point is simply called a **stationary point**.

We now state our main theorem of this section, which shows that for Lipschitz-continuous and Lipschitz-smooth VIs, the regularized primal gap function is weakly-concave, and as such we can compute a stationary point using standard proof techniques.

Theorem 4.4.1 [Mirror potential method convergence]**.**
Let $(\mathcal{X}, \boldsymbol{f})$ be a λ-Lipschitz-continuous and β-Lipschitz-smooth VI, h a 1-strongly-convex kernel function, $\alpha \geq 0$, $\eta \in \left(0, \frac{1}{2(2\beta\alpha\mathrm{diam}(\mathcal{X})^2+1+2\lambda)}\right]$, and $\boldsymbol{x}^{(0)} \in \mathcal{X}$.

Consider the mirror potential algorithm (Algorithm 5) run with the regularized primal gap Ξ_α associated with $(\mathcal{X}, \boldsymbol{f})$, the kernel function h, an arbitrary time horizon $\tau \in \mathbb{N}$, the step size η, and the initial iterate $\boldsymbol{x}^{(0)}$. The following convergence bound to a stationary point of Ξ_α then holds on the outputs $\{\boldsymbol{x}^{(t)}\}_t$:

$$\min_{k=0,1,\ldots,\tau-1} \max_{\boldsymbol{x} \in \mathcal{X}} \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x} \rangle \leq \frac{2\, \Xi_\alpha(\boldsymbol{x}^{(0)})}{\tau}$$

In addition, if $\boldsymbol{x}^{(\tau)}_{\text{best}} \in \arg\min_{\boldsymbol{x}^{(k)}:k=0,\ldots,\tau-1} \max_{\boldsymbol{x} \in \mathcal{X}} \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x} \rangle$, then for some $\tau \in O(1/\varepsilon)$, $\boldsymbol{x}^{(\tau)}_{\text{best}}$ is an ε-stationary point of Ξ_α.

Proof

For convenience, define $\nu \doteq (2\beta\alpha \operatorname{diam}(\mathcal{X})^2 + 1 + 2\lambda)$. By Lemma 4.4.3, Ξ_α is ν-weakly-concave.

Now, for all $k \in \mathbb{N}_+$, by the first-order optimality condition at $\boldsymbol{x}^{(k+1)}$, for all $\boldsymbol{x} \in \mathcal{X}$ and $k \in \mathbb{N}$:

$$\langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}) + \frac{1}{\eta} \nabla h(\boldsymbol{x}^{(k+1)}) - \nabla h(\boldsymbol{x}^{(k)}), \boldsymbol{x} - \boldsymbol{x}^{(k+1)} \rangle \geq 0.$$

Substituting $\boldsymbol{x} = \boldsymbol{x}^{(k)}$ above, for all $k \in \mathbb{N}$,

$$\begin{aligned} \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x}^{(k+1)} \rangle &\geq \frac{1}{\eta} \langle \nabla h(\boldsymbol{x}^{(k+1)}) - \nabla h(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)} \rangle \\ &= \frac{1}{\eta} \left(\operatorname{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k)}) + \operatorname{div}_h(\boldsymbol{x}^{(k)}, \boldsymbol{x}^{(k+1)}) \right) \quad , \end{aligned}$$

where the last line follows from Lemma 4.3.1.

Re-organizing yields for all $k \in \mathbb{N}$,

$$\operatorname{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k)}) \leq \eta \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x}^{(k+1)} \rangle - \underbrace{\operatorname{div}_h(\boldsymbol{x}^{(k)}, \boldsymbol{x}^{(k+1)})}_{\geq 0} \tag{4.40}$$

$$\leq -\eta \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)} \rangle \tag{4.41}$$

Now, by the ν-weak-concavity of Ξ_α, and the 1-strong-convexity of the kernel function h, which implies $\operatorname{div}_h(\boldsymbol{x}, \boldsymbol{y}) \geq 1/2 \|\boldsymbol{x} - \boldsymbol{y}\|^2$, for all $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{X}$, it follows that for all $k \in \mathbb{N}$,

$$
\begin{aligned}
&\Xi_\alpha(\boldsymbol{x}^{(k+1)}) \\
&\leq \Xi_\alpha(\boldsymbol{x}^{(k)}) + \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)} \rangle + \frac{\nu}{2} \left\| \boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)} \right\|^2 \\
&\leq \Xi_\alpha(\boldsymbol{x}^{(k)}) + \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)} \rangle + \nu \mathrm{div}_h(\boldsymbol{x}^{(k+1)}, \boldsymbol{x}^{(k)}) \\
&\leq \Xi_\alpha(\boldsymbol{x}^{(k)}) + (1 - \nu\eta) \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k+1)} - \boldsymbol{x}^{(k)} \rangle && \text{(Equation (4.41))} \\
&\leq \Xi_\alpha(\boldsymbol{x}^{(k)}) + (1 - \nu\eta) \max_{\boldsymbol{x} \in \mathcal{X}} \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x} - \boldsymbol{x}^{(k)} \rangle
\end{aligned}
$$

Unrolling the inequality for $k = 0, 1, \ldots, \tau - 1$, yields:

$$
\Xi_\alpha(\boldsymbol{x}^{(\tau)}) \leq \Xi_\alpha(\boldsymbol{x}^{(0)}) + (1 - \nu\eta) \sum_{k=0}^{\tau-1} \max_{\boldsymbol{x} \in \mathcal{X}} \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x} - \boldsymbol{x}^{(k)} \rangle
$$

Re-organizing the inequality, and dropping the expression $\Xi_\alpha(\boldsymbol{x}^{(\tau)}) \geq 0$, we have:

$$
(1 - \nu\eta) \sum_{k=0}^{\tau-1} \max_{\boldsymbol{x} \in \mathcal{X}} \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x} \rangle \leq \Xi_\alpha(\boldsymbol{x}^{(0)})
$$

Since $\eta \in (0, 1/2\nu]$, we have $(1 - \nu\eta) = (1 - 1/2) \geq 1/2$. Hence,

$$
\sum_{k=0}^{\tau-1} \max_{\boldsymbol{x} \in \mathcal{X}} \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x} \rangle \leq 2\, \Xi_\alpha(\boldsymbol{x}^{(0)})
$$

Multiplying both sides by $\frac{1}{\tau}$, and and applying the generalized means inequality yields:

$$
\frac{1}{\tau} \sum_{k=0}^{\tau-1} \max_{\boldsymbol{x} \in \mathcal{X}} \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x} \rangle \leq 2\, \Xi_\alpha(\boldsymbol{x}^{(0)})
$$

$$
\min_{k=0,1,\ldots,\tau-1} \max_{\boldsymbol{x} \in \mathcal{X}} \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x} \rangle \leq \frac{2\, \Xi_\alpha(\boldsymbol{x}^{(0)})}{\tau}
$$

Finally, we can convert this convergence bound to a finite-time convergence result for any $\varepsilon \geq 0$, by setting $\min_{k=0,1,\ldots,\tau-1} \max_{\boldsymbol{x} \in \mathcal{X}} \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x} \rangle \leq 2\, \Xi_\alpha(\boldsymbol{x}^{(0)})/\tau \leq \varepsilon$ and solving for τ, which implies

$$
\frac{2\, \Xi_\alpha(\boldsymbol{x}^{(0)})}{\varepsilon} \leq \tau
$$

Hence, for some $\tau \in O(1/\varepsilon)$, it holds that $\min_{k=0,1,\ldots,\tau-1} \max_{\boldsymbol{x} \in \mathcal{X}} \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x} \rangle \leq \varepsilon$.

We conclude with a remark on an interpretation of this theorem, before turning to applications of our results.

Remark 4.4.1 [When are stationary points global solutions]**.**
It is well-known that for monotone VIs, stationary points of the regularized primal gap function correspond to strong solutions of the VI (see, for instance, Theorem 3.3. of Fukushima (1992)). As such, the above result implies a strong solution can be computed in polynomial-time via the mirror potential method in monotone, Lipschitz-continuous, and Lipschitz-smooth VIs.

Chapter 5

Walrasian Economies

5.1 Background

A **Walrasian economy** $(m, \mathcal{Z})$ consists of $m \in \mathbb{N}$ **commodities**,[1] with any quantity of each commodity being exchangeable for a quantity of another. The exchange process is governed by a valuation system called **prices**, modeled as a vector $\boldsymbol{p} \in \mathbb{R}_+^m$ s.t. $p_j \geq 0$ is the price of commodity $j \in [m]$.[2] Prices $\boldsymbol{p} \in \mathbb{R}_+^m$ allow the **sale** of $x \in \mathbb{R}_+$ units of any commodity $j \in [m]$ in exchange for the **purchase** of $x\,(p_j/p_k)$ units of any other commodity $k \in [m]$.

For any price system, the economy determines[3] quantities of each commodity that can be bought and sold, with all admissible exchanges being summarized by an **excess demand correspondence** $\mathcal{Z} : \mathbb{R}_+^m \rightrightarrows \mathbb{R}^m$, which for any prices $\boldsymbol{p} \in \mathbb{R}_+^m$ outputs a **set of excess demands** $\mathcal{Z}(\boldsymbol{p}) \subseteq \mathbb{R}^m$ with each excess demand $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$. For any price $\boldsymbol{p} \in \mathbb{R}_+^m$ and excess demand $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$, $z_j(\boldsymbol{p}) \geq 0$ denotes the number of units of commodity $j \in [m]$ **demanded in excess** (i.e., more units of j are bought than sold), while $z_j(\boldsymbol{p}) < 0$ denotes

[1] The "commodity" terminology is used here in the tradition of Arrow and Debreu (1954), and refers to any raw, intermediate, or finished products, as well as labor and services.

[2] The astute reader might notice that in real-world economies the prices of certain commodities can be negative (e.g., prices of oil when storage of excess oil is not possible), and might raise the concern that the model does not account for the possibility of negative prices. However, in these cases the price of the commodity is "negative" only colloquially speaking; rather, the price of an associated commodity is positive. For instance, when the price of oil is negative, companies are no longer selling oil; instead, they are buying a service: the storage of oil. As such, we account for "negative pricing" in the real-world by including as additional commodities, commodities with "negative prices" (e.g., including both oil and the sale of oil as commodities). We will see in Chapter 13, Part III, an example of a Walrasian economy in which commodities are explicitly modeled, and which captures this negative pricing phenomenon.

[3] Here, for narrative simplicity, the economy determining prices can be interpreted as a fictional auctioneer announcing prices in the tradition of Walras (Walras, 1896).

the number of units of commodity $j \in [m]$ **supplied in excess** (i.e., more units of it are sold than bought). If $\mathcal{Z}$ is singleton-valued, then we will for convenience represent $\mathcal{Z}$ as a function and denote it $\boldsymbol{z}$.

A price vector $\boldsymbol{p} \in \mathbb{R}_+^m$ is **feasible** if there exists a $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$ s.t. for all commodities $j \in [m]$, $z_j(\boldsymbol{p}) \leq 0$. Similarly, a price vector $\boldsymbol{p} \in \mathbb{R}_+^m$ satisfies **Walras' law** if there exists a $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$ s.t. $\boldsymbol{p} \cdot \boldsymbol{z}(\boldsymbol{p}) = 0$.

The canonical solution concept for Walrasian equilibria is the Walrasian equilibrium (Walras, 1896). In the sequel, we will introduce algorithms with polynomial-time convergence guarantees to a Walrasian equilibrium. To analyze their convergence will use a computationally relevant generalization of Walrasian equilibrium, namely the approximate Walrasian equilibrium, to account for the bounded accuracy of computational methods.

Definition 5.1.1 [Approximate Walrasian Equilibrium]**.**
Given an approximation parameter $\varepsilon \geq 0$, a price vector $\boldsymbol{p}^* \in [0,1]^m$ is said to be an ε-**Walrasian** (or ε-**competitive**) **equilibrium** of a Walrasian economy $(m, \mathcal{Z})$ if there exists an excess demand $\boldsymbol{z}(\boldsymbol{p}^*) \in \mathcal{Z}(\boldsymbol{p}^*)$ s.t.

$$
\begin{aligned}
&(\varepsilon-\text{Feasility}) \quad && \text{For all commodities } j \in [m],\ z_j(\boldsymbol{p}^*) \leq \varepsilon \\
&(\varepsilon\text{-Walras' law}) \quad && -\varepsilon \leq \boldsymbol{p}^* \cdot \boldsymbol{z}(\boldsymbol{p}^*) \leq \varepsilon
\end{aligned}
$$

We denote the set of ε-Walrasian equilibria of any Walrasian economy $(m, \mathcal{Z})$ by $\mathcal{WE}_\varepsilon(m, \mathcal{Z})$.

A 0-Walrasian equilibrium is simply called a **Walrasian equilibrium**, in which case we denote the set of Walrasian equilibria $\mathcal{WE}(m, \mathcal{Z})$.

Seen otherwise, a Walrasian equilibrium $\boldsymbol{p}^* \in \mathbb{R}_+^m$ is a price vector s.t. for all commodities $j \in [m]$, $p_j^* > 0 \implies z_j(\boldsymbol{p}^*) = 0$ and $p_j^* = 0 \implies z_j(\boldsymbol{p}^*) \leq 0$. Intuitively, a Walrasian equilibrium is a price vector which ensures that the exchange of any commodity with another can be implemented. On the one hand, if the price of a commodity $j \in [m]$ is strictly positive, then the exchange system dictates that j can be exchanged for a strictly positive quantity of some other commodity $k \in [m]$; in other words, at a Walrasian equilibrium commodity j will always find a buyer since its excess demand is zero. On the other hand, if the price of commodity j is zero, then the price system dictates that the commodity j cannot be exchanged for any other commodity; in other words, at a Walrasian equilibrium the commodity might not find a buyer.

5.2 Walrasian Economies and Variational Inequalities

With definitions in order, we now present the fundamental relationship between Walrasian economies and VIs. We will use this relationship to first establish the existence of a Walrasian equilibrium in continuous balanced economies, and then introduce efficient algorithms for the computation of a Walrasian equilibrium in Lipschitz-continuous balanced economies.

5.2.1 Walrasian Economies and Complementarity Problems

The following theorem, due to Dafermos (1990), is to the best of our knowledge the first result exposing the connection between VIs and Walrasian equilibria (see, Nagurney (2009) for additional references). It states that the problem of computing a Walrasian equilibrium is equivalent to the problem of computing a strong solution of a VI whose set of constraints is given by the positive orthant (a class of VIs known as **complementarity problems** (Cottle and Dantzig, 1968)).

Theorem 5.2.1 [Walrasian economies as Complementarity Problems]**.**
The set of Walrasian equilibria of any Walrasian economy $(m, \mathcal{Z})$ is equal to the set of strong solutions of the VI $(\mathbb{R}^m_+, -\mathcal{Z})$, i.e., $\mathcal{WE}(m, \mathcal{Z}) = \mathcal{SVI}(\mathbb{R}^m_+, -\mathcal{Z})$.

Proof of Theorem 5.2.1

$(\implies)$ Let $\boldsymbol{p}^* \in \mathcal{WE}(m, \mathcal{Z})$ be a Walrasian equilibrium. Then, for some $\boldsymbol{z}(\boldsymbol{p}^*) \in \mathcal{Z}(\boldsymbol{p}^*)$, we have:

$$\begin{aligned}
&\langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} - \boldsymbol{p}^* \rangle && \forall \boldsymbol{p} \in \mathbb{R}^m_+ \\
&= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} \rangle - \underbrace{\langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle}_{=0} && \forall \boldsymbol{p} \in \mathbb{R}^m_+ \\
&= \underbrace{\langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} \rangle}_{\leq 0} && \forall \boldsymbol{p} \in \mathbb{R}^m_+ \\
&\leq 0
\end{aligned}$$

where the last line follows from the feasibility of $\boldsymbol{z}(\boldsymbol{p}^*)$, i.e., $\boldsymbol{z}(\boldsymbol{p}^*) \leq 0$, and the positivity of $\boldsymbol{p}$.

$(\impliedby)$ Let $\boldsymbol{p}^* \in \mathcal{SVI}(\mathbb{R}^m_+, -\mathcal{Z})$. Then, for some $\boldsymbol{z}(\boldsymbol{p}^*) \in \mathcal{Z}(\boldsymbol{p})$, we have:

$$0 \geq \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} - \boldsymbol{p}^* \rangle \qquad \forall \boldsymbol{p} \in \mathbb{R}^m_+$$

Substituting $\boldsymbol{p} \doteq \boldsymbol{p}^* + \boldsymbol{j}_j$, we have:

$$\begin{aligned}
0 &\geq \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* + \boldsymbol{j}_j - \boldsymbol{p}^* \rangle \\
&= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{j}_j \rangle \\
&\geq z_j(\boldsymbol{p}^*) && \forall j \in [m]
\end{aligned}$$

That is, $\boldsymbol{p}^*$ is feasible.

Similarly, substituting in $\boldsymbol{p} \doteq \boldsymbol{0}$ and $\boldsymbol{p} \doteq 2\boldsymbol{p}^*$, we have:

$$0 \leq \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle$$

and

$$0 \geq \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle$$

That is, $\boldsymbol{p}^*$ satisfies weak Walras' law. Hence, $\boldsymbol{p}^*$ is a Walrasian equilibrium.

5.2.2 Balanced Economies and Variational Inequalities

While Theorem 5.2.1 shows that we can approach any Walrasian equilibrium computation problem as a strong VI computation problem, because the domain of prices is unbounded (i.e., $\mathbb{R}_+^m$), to obtain existence and convergence results, we have to restrict the class of Walrasian economies we study. To this end, we introduce two important classes of Walrasian economies. The first of these classes is balanced economies.

Definition 5.2.1 [Balanced economies]**.**
A **balanced economy** is a Walrasian economy $(m, \mathcal{Z})$ whose excess demand correspondence satisfies:

$$\text{(Homogeneity of degree 0)} \qquad \text{For all } \lambda > 0,\ \mathcal{Z}(\lambda \boldsymbol{p}) = \mathcal{Z}(\boldsymbol{p})$$

$$\text{(Weak Walras' law)} \qquad \text{For all } \boldsymbol{p} \in \mathbb{R}_+^m \text{ and } \boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p}),\ \boldsymbol{p} \cdot \boldsymbol{z}(\boldsymbol{p}) \leq 0$$

Intuitively, homogeneity requires that prices have a meaning only relative to other prices, and have no absolute meaning of their own (i.e., if all prices are scaled by the same amount, the excess demand is unchanged); weak Walras' law requires budget-balance (i.e., at all prices, the total value of what is being demanded cannot exceed the total value of what is being supplied). While homogeneity of degree 0 is a standard assumption, weak Walras' law is significantly weaker than standard assumptions previously considered in the literature (see, for instance Arrow and Hurwicz (1958) and Debreu (1974)), and is satisfied by Arrow-Debreu economies (Arrow and Debreu, 1954) (see, Chapter 10 for additional details).

We now provide a novel characterization of Walrasian equilibrium prices in balanced economies as a VI over $[0, 1]^m$ rather than $\mathbb{R}_+^m$, which will allow us to obtain polynomial-time algorithms for the computation of Walrasian equilibrium, as the computational guarantees of our algorithms for VIs depend on the diameter of the constraint space of the VIs. In particular, we will now show that the set of Walrasian equilibria of any balanced economy can be restated as the set of strong solutions of a modified VI $([0, 1]^m, -\mathcal{Z})$ where the constraint space is $[0,1]^m$: i.e.,

$$\text{Find } \boldsymbol{p}^* \in [0,1]^m \text{ and } \boldsymbol{z}(\boldsymbol{p}^*) \in \mathcal{Z}(\boldsymbol{p}^*) \text{ s.t. for all } \boldsymbol{p} \in [0,1]^m, \quad \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} - \boldsymbol{p}^* \rangle \leq 0 \tag{5.1}$$

Theorem 5.2.2 [Balanced economies as VIs]**.**

For any balanced economy $(m, \mathcal{Z})$, the set of Walrasian equilibria is equal to the strictly positive cone generated by the strong solutions of the continuous VI $([0,1]^m, -\mathcal{Z})$, i.e., $\mathcal{WE}(m, \mathcal{Z}) = \bigcup_{\lambda \geq 1} \lambda \, \mathcal{SVI}([0,1]^m, -\mathcal{Z})$.

Proof of Theorem 5.2.2

($\implies$) Let $\boldsymbol{p}^* \in \mathcal{WE}(m, \mathcal{Z})$ be a Walrasian equilibrium. Let $\alpha \doteq \frac{1}{\max\{1, \|\boldsymbol{p}^*\|_\infty\}}$ so that $\alpha \boldsymbol{p}^* \in [0,1]^m$. Now, for some $\boldsymbol{z}(\alpha \boldsymbol{p}^*) \in \mathcal{Z}(\alpha \boldsymbol{p}^*)$, we have:

$$\begin{aligned}
&\langle -\boldsymbol{z}(\alpha \boldsymbol{p}^*), \alpha \boldsymbol{p}^* - \boldsymbol{p} \rangle && \forall \boldsymbol{p} \in [0,1]^m \\
&= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} - \alpha \boldsymbol{p}^* \rangle && \forall \boldsymbol{p} \in [0,1]^m && \text{(Homogeneity of } \boldsymbol{z}) \\
&= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} \rangle - \alpha \underbrace{\langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle}_{=0} && \forall \boldsymbol{p} \in [0,1]^m \\
&= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} \rangle && \forall \boldsymbol{p} \in [0,1]^m \\
&\leq 0
\end{aligned}$$

where the penultimate line follows from Walras' law holding at a Walrasian equilibrium, and the last line follows from the feasibility of $\boldsymbol{z}(\boldsymbol{p}^*)$, i.e., $\boldsymbol{z}(\boldsymbol{p}^*) \leq \boldsymbol{0}$ and the positivity of $\boldsymbol{p}$. Hence, $\alpha \boldsymbol{p}^*$ is a strong solution of the VI $([0,1]^m, -\mathcal{Z})$, which means that $\boldsymbol{p}^* \in {}^1\!/\!_\alpha \, \mathcal{SVI}([0,1]^m, -\mathcal{Z})$.

Now, notice that by the homogeneity of the excess demand in balanced economies, since for all $\lambda > 0$, $\mathcal{Z}(\lambda \boldsymbol{p}^*) = \mathcal{Z}(\boldsymbol{p}^*)$, if $\boldsymbol{p}^*$ is a Walrasian equilibrium, then so is $\lambda \boldsymbol{p}^*$. Hence, α takes values in $(0,1]$, implying ${}^1\!/\!_\alpha \in [1, \infty)$. As such, we must have $\mathcal{WE}(m, \mathcal{Z}) \subseteq \bigcup_{\lambda \geq 1} \lambda \, \mathcal{SVI}([0,1]^m, -\mathcal{Z})$.

($\Longleftarrow$) Let $\boldsymbol{p}^* \in \mathcal{SVI}([0,1]^m, -\mathcal{Z})$ and $\lambda \geq 1$. Then, for some $\boldsymbol{z}(\boldsymbol{p}^*) \in \mathcal{Z}(\boldsymbol{p}^*)$, we have:

$$
\begin{aligned}
0 &\geq \langle -\boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* - \boldsymbol{p} \rangle && \forall \boldsymbol{p} \in [0,1]^m \\
&= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} - \boldsymbol{p}^* \rangle && \forall \boldsymbol{p} \in [0,1]^m \\
&= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} \rangle - \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle && \forall \boldsymbol{p} \in [0,1]^m \qquad (5.2)
\end{aligned}
$$

Plugging $\boldsymbol{p} = \mathbf{0}_m$ into Equation (5.2), we then have:

$$
\begin{aligned}
0 &\geq \underbrace{\langle \boldsymbol{z}(\boldsymbol{p}^*), \mathbf{0}_m \rangle}_{=0} - \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle \\
0 &\geq - \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle \\
0 &\leq \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle \\
0 &\leq \langle \boldsymbol{z}(\lambda \boldsymbol{p}^*), \boldsymbol{p}^* \rangle && \text{(Homogeneity of } \boldsymbol{z}\text{)} \\
0 &\leq \langle \boldsymbol{z}(\lambda \boldsymbol{p}^*), \lambda \boldsymbol{p}^* \rangle
\end{aligned}
$$

Further, since $(m, \mathcal{Z})$ is balanced, we have $\lambda \boldsymbol{p}^* \cdot \boldsymbol{z}(\lambda \boldsymbol{p}^*) = \boldsymbol{p}^* \cdot \boldsymbol{z}(\boldsymbol{p}^*) \leq 0$, hence, combining it with the above inequality, we must have $\lambda \boldsymbol{p}^* \cdot \boldsymbol{z}(\lambda \boldsymbol{p}^*) = 0$, meaning that $\lambda \boldsymbol{p}^*$ satisfies Walras' law.

In addition, continuing from Equation (5.2) again, we have:

$$
\begin{aligned}
0 &\geq \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} \rangle - \underbrace{\langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle}_{\leq 0} \\
&\geq \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} \rangle && \forall \boldsymbol{p} \in [0,1]^m \\
&\geq \langle \boldsymbol{z}(\lambda \boldsymbol{p}^*), \boldsymbol{p} \rangle && \forall \boldsymbol{p} \in [0,1]^m \qquad (5.3)
\end{aligned}
$$

where the penultimate line follows from the fact that balanced economies satisfy weak Walras' law, and the last line from the homogeneity of degree 0 of the excess demand.

Now, plugging $\boldsymbol{p} = \boldsymbol{j}_j$ into Equation (5.3), we have:

$$
\begin{aligned}
0 &\geq \langle \boldsymbol{z}(\lambda \boldsymbol{p}^*), \boldsymbol{j}_j \rangle && \forall j \in [m] \\
&\geq z_j(\lambda \boldsymbol{p}^*) && \forall j \in [m]
\end{aligned}
$$

That is, $\lambda \boldsymbol{p}^*$ is feasible. Putting it all together, $\lambda \boldsymbol{p}^*$ must be a Walrasian equilibrium. As such, $\bigcup_{\lambda \geq 1} \lambda\, \mathcal{SVI}([0,1]^m, -\mathcal{Z}) \subseteq \mathcal{WE}(m, \mathcal{Z})$.

In the sequel, we will make use of the following lemma, which states that for any balanced economy $(m, \mathcal{Z})$, any approximate strong solution of the VI $([0,1]^m, -\mathcal{Z})$ is an approximate Walrasian equilibrium of $(m, \mathcal{Z})$.

Lemma 5.2.1 [ε-strong solution and ε-Walrasian equilibrium]**.**

For any balanced economy $(m, \mathcal{Z})$, any ε-strong solution of the VI $([0,1]^m, -\mathcal{Z})$ is an ε-Walrasian equilibrium of $(m, \mathcal{Z})$.

Proof of Lemma 5.2.1

(ε-strong solution $\implies$ ε-Walrasian equilibrium) For any $\varepsilon \geq 0$, let $\boldsymbol{p}^* \in \mathcal{SVI}_\varepsilon([0,1]^m, -\mathcal{Z})$. Then, for some $\boldsymbol{z}(\boldsymbol{p}^*) \in \mathcal{Z}(\boldsymbol{p}^*)$, we have:

$$\begin{aligned} \varepsilon &\geq \langle -\boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* - \boldsymbol{p} \rangle && \forall \boldsymbol{p} \in [0,1]^m \\ &= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} - \boldsymbol{p}^* \rangle && \forall \boldsymbol{p} \in [0,1]^m \\ &= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} \rangle - \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle && \forall \boldsymbol{p} \in [0,1]^m \end{aligned} \tag{5.4}$$

Plugging $\boldsymbol{p} = \mathbf{0}_m$ into Equation (5.2), we then have:

$$\begin{aligned} \varepsilon &\geq \underbrace{\langle \boldsymbol{z}(\boldsymbol{p}^*), \mathbf{0}_m \rangle}_{=0} - \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle \\ \varepsilon &\geq - \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle \\ -\varepsilon &\leq \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle \end{aligned} \tag{5.5}$$

Further, since $(m, \mathcal{Z})$ is balanced, it follows that $\boldsymbol{p}^* \cdot \boldsymbol{z}(\boldsymbol{p}^*) \leq 0 \leq \varepsilon$. Combining this conclusion with Equation (5.5), we see that $\boldsymbol{p}^*$ satisfies ε-Walras' law.

In addition, continuing from Equation (5.4) again, we have:

$$\begin{aligned} \varepsilon &\geq \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} \rangle - \underbrace{\langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle}_{\leq 0} \\ &\geq \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} \rangle && \forall \boldsymbol{p} \in [0,1]^m \end{aligned} \tag{5.6}$$

where the last line follows from the fact that balanced economies satisfy weak Walras' law.

Now, plugging $\boldsymbol{p} = \boldsymbol{j}_j$ into Equation (5.6), we have:

$$\begin{aligned} \varepsilon &\geq \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{j}_j \rangle & \forall j \in [m] \\ &\geq z_j(\boldsymbol{p}^*) & \forall j \in [m] \end{aligned} .$$

That is, $\boldsymbol{p}^*$ is ε-feasible. Putting it all together, $\boldsymbol{p}^*$ must be an ε-Walrasian equilibrium.

5.2.3 Competitive Economies and Continuous Variational Inequalities

We now turn our attention to proving the existence of Walrasian equilibrium. In balanced economies, under the assumption that the excess demand correspondence $\mathcal{Z}$ is upper hemicontinuous, non-empty-, compact-, and convex-valued, the existence of a Walrasian equilibrium $\boldsymbol{p}^* \in [0,1]^m$ follows as a corollary of the existence of strong solutions to continuous VIs (Theorem 4.1.1). Unfortunately, this Walrasian equilibrium can be trivial, i.e., $\boldsymbol{p}^* = \mathbf{0}_m$ is an equilibrium. To prove the existence of a non-trivial Walrasian equilibrium, we have to restrict our attention to a subset of balanced Walrasian economies. We choose to study a canonical subset (Debreu, 1974; Sonnenschein, 1972), which we call **competitive economies**.

Definition 5.2.2 [Competitive economy]**.**

A **competitive economy** is a Walrasian economy $(m, \mathcal{Z})$ whose excess demand correspondence satisfies:

$$\begin{aligned} &\text{(Homogeneity of degree 0)} && \text{For all } \lambda > 0,\ \mathcal{Z}(\lambda \boldsymbol{p}) = \mathcal{Z}(\boldsymbol{p}) \\ &\text{(Weak Walras' law)} && \text{For all } \boldsymbol{p} \in \mathbb{R}^m_+ \text{ and } \boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p}),\ \boldsymbol{p} \cdot \boldsymbol{z}(\boldsymbol{p}) \leq 0 \\ &\text{(Non-Satiation)} && \text{For all } \boldsymbol{p} \in \mathbb{R}^m_+ \text{ and } \boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p}),\ \boldsymbol{z}(\boldsymbol{p}) \leq \mathbf{0}_m \text{ implies } \boldsymbol{p} \cdot \boldsymbol{z}(\boldsymbol{p}) = 0 \end{aligned}$$

That is, a competitive economy is a balanced economy with the additional requirement that when the excess demand is feasible, Walras' law holds. Intuitively, this non-satiation condition requires that whenever all commodities are supplied in excess, it must be that the economy has exhausted its purchasing power. As such, the excess demand is non-satiated, in the sense that the economy cannot demand more of any commodity, not because it is not supplied in sufficient quantity, but rather because it cannot afford it.

In competitive economies, an alternative VI characterization of Walrasian equilibrium holds over the constraint space Δ_m rather than $[0,1]^m$, which is more suitable for proving existence.

Theorem 5.2.3 [Competitive economies as VIs]**.**
For any competitive economy $(m, \mathcal{Z})$, the set of Walrasian equilibria is equal to the strictly positive cone generated by the strong solutions of the continuous VI $(\Delta_m, -\mathcal{Z})$, i.e., $\mathcal{WE}(m, \mathcal{Z}) = \bigcup_{\lambda \geq 1} \lambda\, \mathcal{SVI}(\Delta_m, -\mathcal{Z})$.

Proof of Theorem 5.2.3

($\implies$) Let $\boldsymbol{p}^* \in \mathcal{WE}(m, \mathcal{Z})$ be a Walrasian equilibrium. Let $\alpha \doteq \frac{1}{\|\boldsymbol{p}^*\|_1}$. Then, we have $\alpha \boldsymbol{p}^* \in \Delta_m$. Further, for some $\boldsymbol{z}(\alpha \boldsymbol{p}^*) \in \mathcal{Z}(\alpha \boldsymbol{p}^*)$, we have:

$$
\begin{aligned}
&\langle -\boldsymbol{z}(\alpha \boldsymbol{p}^*), \alpha \boldsymbol{p}^* - \boldsymbol{p} \rangle && \forall \boldsymbol{p} \in \Delta_m \\
&= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} - \alpha \boldsymbol{p}^* \rangle && \forall \boldsymbol{p} \in \Delta_m && \text{(Homogeneity of } \boldsymbol{z}\text{)} \\
&= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} \rangle - \alpha \underbrace{\langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* \rangle}_{=0} && \forall \boldsymbol{p} \in \Delta_m \\
&= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} \rangle && \forall \boldsymbol{p} \in \Delta_m \\
&\leq 0
\end{aligned}
$$

where the penultimate line follows from Walras' law holding at a Walrasian equilibrium, and the last line follows from the feasibility of $\boldsymbol{z}(\boldsymbol{p}^*)$, i.e., $\boldsymbol{z}(\boldsymbol{p}^*) \leq \boldsymbol{0}$, and the positivity of $\boldsymbol{p}$. Hence, $\alpha \boldsymbol{p}^*$ is a strong solution of the VI $(\Delta_m, -\mathcal{Z})$, which means that $\boldsymbol{p}^* \in \frac{1}{\alpha} \mathcal{SVI}(\Delta_m, -\mathcal{Z})$.
Now, notice that by homogeneity of the excess demand in competitive economies since for all $\lambda > 0$, $\mathcal{Z}(\lambda \boldsymbol{p}^*) = \mathcal{Z}(\boldsymbol{p}^*)$, if $\boldsymbol{p}^*$ is a Walrasian equilibrium, then so is $\lambda \boldsymbol{p}^*$.

Hence, α takes values in $(0,1]$, implying $\frac{1}{\alpha} \in [1,\infty)$, and as such we must have $\mathcal{WE}(m,\mathcal{Z}) \subseteq \bigcup_{\lambda\geq 1} \lambda\, \mathcal{SVI}(\Delta_m, -\mathcal{Z})$.

($\Longleftarrow$) Let $\boldsymbol{p}^* \in \mathcal{SVI}(\Delta_m, -\mathcal{Z})$ and $\lambda \geq 1$. Then, for some $\boldsymbol{z}(\boldsymbol{p}^*) \in \mathcal{Z}(\boldsymbol{p}^*)$, we have:

$$
\begin{aligned}
0 &\geq \langle -\boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^* - \boldsymbol{p}\rangle && \forall \boldsymbol{p} \in \Delta_m \\
&= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p} - \boldsymbol{p}^*\rangle && \forall \boldsymbol{p} \in \Delta_m \\
&= \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}\rangle - \underbrace{\langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}^*\rangle}_{\leq 0} \\
&\geq \langle \boldsymbol{z}(\boldsymbol{p}^*), \boldsymbol{p}\rangle && \forall \boldsymbol{p} \in \Delta_m \\
&\geq \langle \boldsymbol{z}(\lambda\boldsymbol{p}^*), \boldsymbol{p}\rangle && \forall \boldsymbol{p} \in \Delta_m
\end{aligned}
$$

where the penultimate line follows from the fact that competitive economies satisfy weak Walras' law, and the last line from homogeneity of degree 0 of the excess demand.

Now, plugging $\boldsymbol{p} = \boldsymbol{j}_j$ for all $j \in [m]$ in the above, we have:

$$
\begin{aligned}
0 &\geq \langle \boldsymbol{z}(\lambda\boldsymbol{p}^*), \boldsymbol{j}_j\rangle && \forall j \in [m] \\
&\geq z_j(\lambda\boldsymbol{p}^*) && \forall j \in [m] \;.
\end{aligned}
$$

That is, $\lambda\boldsymbol{p}^*$ is feasible. Now by non-satiation, since $\boldsymbol{z}(\lambda\boldsymbol{p}^*) \leq \boldsymbol{0}_m$, we must have $\lambda\boldsymbol{p}^* \cdot \boldsymbol{z}(\lambda\boldsymbol{p}^*) \geq 0$. As by weak Walras' law $\lambda\boldsymbol{p}^* \cdot \boldsymbol{z}(\lambda\boldsymbol{p}^*) \leq 0$, we must have $\lambda\boldsymbol{p}^* \cdot \boldsymbol{z}(\lambda\boldsymbol{p}^*) = 0$, meaning that $\lambda\boldsymbol{p}^*$ satisfies Walras' law. Putting it all together, $\lambda\boldsymbol{p}^*$ must be a Walrasian equilibrium. As such we must have $\bigcup_{\lambda\geq 1} \lambda\, \mathcal{SVI}(\Delta_m, -\mathcal{Z}) \subseteq \mathcal{WE}(m,\mathcal{Z})$

To prove existence, it will be necessary to make assumptions on the continuity of the excess demand, which necessitates the definition of continuous economies. We note that in the following definition we assume upper hemicontinuity only on Δ_m, since in competitive, and more generally balanced, economies it is too restrictive to assume that the excess demand $\mathcal{Z}$ is upper hemicontinuous on $\mathbb{R}^m_+$, since any correspondence which is homogeneous of degree 0 and continuous in the entirety of its domain is constant.

Intuitively, continuous economies are those economies in which changes in the proportions of prices lead to well-behaved changes in excess demands.

Definition 5.2.3 [Continuous economies]**.**

A **continuous economy** is a Walrasian economy $(m, \mathcal{Z})$ whose excess demand correspondence $\mathcal{Z}$ is upper hemicontinuous on Δ_m, non-empty-, compact-, and convex-valued.

Remark 5.2.1 [Continuity of excess demand]**.**

In more stylized applications (see, for instance, Chapter 6 or Chapter 10), the excess demand correspondence is in general defined so as to be continuous only on the interior of the unit simplex, i.e., $\mathrm{int}(\Delta_m)$, as the excess demand for a commodity can be infinite if its price is 0. However, this issue in these stylized models only arises from a modeling choice, which allows the demand of commodities to exceed the total amount of the commodity that can be ever supplied. It is indeed possible to restrict the excess demand for a commodity to be bounded by the total amount of the commodity that can be ever supplied in the economy, without modifying the set of Walrasian equilibria. Arrow and Debreu (1954) take exactly this approach in Section 3 of their paper when proving their seminal Walrasian equilibrium existence result, and it is also the approach we will take in Chapter 10 to prove convergence of price adjustment processes in Arrow-Debreu economies. This restriction is also realistic from an economic perspective, since it is not possible for the economy to consume more of a commodity that can exist, and resources in the real-world are indeed scarce. Indeed, otherwise there would be no use for the economic sciences: the science of resource allocation under scarcity.

With the above theorem in hand, we can leverage the fact that a strong solution is guaranteed to exist in continuous VIs (Theorem 4.1.1) to establish the existence of a Walrasian equilibrium in continuous competitive economies.

Theorem 5.2.4.

The set of Walrasian equilibria of any continuous competitive economy $(m, \mathcal{Z})$ is non-empty, i.e., $\mathcal{WE}(m, \mathcal{Z}) \neq \emptyset$.

Proof of Theorem 5.2.4

By Theorem 5.2.3, we know that the set of strong solutions $\mathcal{SVI}(\Delta_m, -\mathcal{Z})$ of the VI $(\Delta_m, -\mathcal{Z})$ is a subset of the set of Walrasian equilibria of a competitive economy. Now, notice that for a continuous competitive economy $(m, \mathcal{Z})$, the corresponding VI $(\Delta_m, -\mathcal{Z})$ is continuous. Hence, by Theorem 4.1.1, a strong solution to $(\Delta_m, -\mathcal{Z})$ is guaranteed to exist, which in turn implies the existence of a Walrasian equilibrium in continuous competitive economies.

5.3 Algorithms for Walrasian Equilibrium

5.3.1 Computational Model

To analyze the computational properties of algorithms for computing Walrasian equilibrium, we will take two approaches. For balanced economies, we will develop algorithms with polynomial-time global convergence guarantees. For general Walrasian economies, we do not expect such algorithms to exist, as by Theorem 5.2.1 the computation of a Walrasian equilibrium is equivalent the PPAD-hard problem of solving a complementarity problem (see, for instance, IEOR (2011)). As such, we will instead define a merit function (i.e., a function whose set of minimizers coincide with the set of Walrasian equilibria), and provide polynomial-time convergence guarantees to approximate stationary points of this merit function (as defined in Chapter 4).

In the rest of this chapter, we will assume that the excess demand correspondence is singleton-valued unless otherwise noted. Similar to Chapter 4, we will consider two classes of methods to compute a Walrasian equilibrium, (first-order) price-adjustment processes and second-order price-adjustment processes, both of which belong to the class of kth order price adjustment processes.

Definition 5.3.1 [kth-order price-adjustment process]**.**
Given some $k \in \mathbb{N}_{++}$, and a Walrasian economy $(m, \mathcal{Z})$ for which the derivatives $\{\nabla^j \boldsymbol{z}\}_{j=1}^{k-1}$ are well defined, and an initial iterate $\boldsymbol{p}^{(0)} \in \mathbb{R}^m_+$, a k**th-order price adjustment process** $\boldsymbol{\pi}$

consists of an update function that generates the sequence of iterates $\{\boldsymbol{p}^{(t)}\}_t$ given by: for all $t = 0, 1, \ldots,$

$$\boldsymbol{x}^{(t+1)} \doteq \boldsymbol{\pi}\left(\bigcup_{i=0}^{t}\left(\boldsymbol{p}^{(i)}, \{\nabla^j \boldsymbol{z}(\boldsymbol{p}^{(i)})\}_{j=0}^{k-1}\right)\right)$$

The computational complexity results in this chapter will rely on the following computational model, which has been broadly adopted in the literature (see, for instance, Papadimitriou and Yannakakis (2010)).

Definition 5.3.2 [Walrasian Computational Model]**.**
Given a Walrasian economy $(m, \mathcal{Z})$ and a kth-order price adjustment process $\boldsymbol{\pi}$, the computational complexity of $\boldsymbol{\pi}$ is measured in term of the number of evaluations of the the functions $\boldsymbol{z}, \nabla \boldsymbol{z}, \ldots, \nabla^k \boldsymbol{z}$.

5.3.2 Related Works

We review here some of the relevant computer science literature, on price-adjustment processes for Walrasian equilibrium computation. This literature has in part been motivated by applications of algorithms such as *tâtonnement* to load balancing over networks (Jain et al., 2013) and to pricing of transactions on crypotocurrency blockchains (Leonardos et al., 2021; Liu et al., 2022; Reijsbergen et al., 2021). A detailed inquiry into the computational properties of Walrasian equilibria was initiated by Devanur et al. (2008), who studied a special case of the Arrow-Debreu competitive economy known as the **Fisher market** (Brainard et al., 2000). This model, for which Irving Fisher computed equilibrium prices using a hydraulic machine in the 1890s, is essentially the Arrow-Debreu model of a competitive economy, but there are no firms, and buyers are endowed with only one type of commodity—an artificial currency (Brainard et al., 2000; Nisan and Roughgarden, 2007). Devanur et al. (2002) exploited a connection first made by Eisenberg (1961) between the Eisenberg-Gale program and competitive equilibrium to solve Fisher markets assuming buyers with linear utility functions, thereby providing a (centralized) polynomial-time algorithm for equilibrium computation in these markets (Devanur et al., 2002; Devanur et al., 2008). Their work was built upon by Jain et al. (2005), who extended the Eisenberg-Gale program to all Fisher markets in which buyers have continuous, quasi-concave, and homogeneous utility functions, and proved that the equilibrium of Fisher markets with such buyers can be computed in polynomial time by interior point methods.

Concurrent with this line of work on computing competitive equilibrium using centralized methods, a line of work on devising and proving convergence guarantees for price-adjustment processes (i.e., iterative algorithms that update prices according to a predetermined update rule) developed. This literature has focused on devising *natural* price-adjustment processes, like *tâtonnement*, which might explain or imitate the movement of prices in real-world markets. In addition to imitating the law of supply and demand, *tâtonnement* has been observed to replicate the movement of prices in lab experiments, where participants are given endowments and asked to trade with one another (Gillen et al., 2020). Perhaps more importantly, the main premise for research on the computational properties of competitive equilibrium in computer science is that for competitive equilibrium to be justified, not only should it be backed by a natural price-adjustment process as economists have long argued, but it should also be computationally efficient (Nisan and Roughgarden, 2007).

The first result on this question is due to Codenotti et al. (2005), who introduced a discrete-time version of *tâtonnement*, and showed that in exchange economies that satisfy **weak gross substitutes (WGS)**, the *tâtonnement* process converges to an approximate competitive equilibrium in a number of steps which is polynomial in the approximation factor and size of the problem. Unfortunately, soon after this positive result appeared, Papadimitriou and Yannakakis (2010) showed that it is impossible for a price-adjustment process based on the excess demand function to converge to a competitive equilibrium in polynomial time in general competitive economies, ruling out the possibility of Smale's process (and many others) justifying the notion of competitive equilibrium in all competitive economies. Nevertheless, further study of the convergence of price-adjustment processes such as *tâtonnement* under stronger assumptions, or in simpler models than full-blown Arrow-Debreu competitive economies, remains worthwhile, as these processes are being deployed in practice (Jain et al., 2013; Leonardos et al., 2021; Liu et al., 2022; Reijsbergen et al., 2021).

Following Codenotti et al.'s (2005) initial analysis of *tâtonnement* in competitive economies that satisfy WGS, Garg and Kapoor (2004) introduced an auction algorithm that also converges in polynomial time for linear exchange economies. More recently, Bei et al. (2015) established faster convergence bounds for *tâtonnement* in WGS exchange economies.

Another line of work considers price-adjustment processes in variants of Fisher markets. Cole and Fleischer (2008) analyzed *tâtonnement* in a real-world-like model satisfying WGS called the ongoing market model. In this model, *tâtonnement* once-again converges in polynomial-time (Cole and Fleischer, 2008; Cole et al., 2010), and it has the advantage that it can be seen as an abstraction for market processes. Cole and Fleischer's results were later extended by Cheung et al. (2012) to ongoing markets with **weak gross complements** (i.e., the excess demand of any commodity weakly increases if the price of any other commodity weakly decreases, fixing all other prices) and ongoing markets with a mix of WGC and WGS commodities. The ongoing market model studied in these two papers contains as a special case the Fisher market; however, Cole and Fleischer (2008) assume bounded own-price elasticity of Marshallian demand, and bounded income elasticity of Marshallian demand, while Cheung et al. (2012) assume, in addition to Cole and Fleischer's assumptions, bounded adversarial market elasticity, which can be seen as a variant of bounded cross-price elasticity of Marshallian demand, from below. With these assumptions, these results cover Fisher markets with only a small range of the well known CES utilities, namely CES Fisher markets with $\rho \in [0, 1)$ and WGC Fisher markets with $\rho \in (-1, 0]$.

Cheung et al. (2013) built on this work by establishing the convergence of *tâtonnement* in polynomial time in nested CES Fisher markets, excluding the limiting cases of linear and Leontief markets, but nonetheless extending polynomial-time convergence guarantees for *tâtonnement* to Leontief Fisher markets as well. More recently, Cheung and Cole (2018) showed that Cheung et al.'s [2013] result extends to an asynchronous version of *tâtonnement*, in which good prices are updated during different time periods. In a similar vein, Cheung et al. (2019) analyzed *tâtonnement* in online Fisher markets, determining that *tâtonnement* tracks competitive equilibrium prices closely provided the market changes slowly.

Another price-adjustment process that has been shown to converge to Walrasian equilibrium in Fisher markets is **proportional response dynamics**, first introduced by Wu and Zhang (2007) for linear utilities; then expanded upon and shown to converge by Zhang (2011) for all CES utilities; and very recently shown to converge in Arrow-Debreu exchange economies with linear and CES ($\rho \in [0, 1)$) utilities by Brânzei et al. (2021). The study of the proportional response process was proven fundamental when Cheung et al. (2013) noticed its relationship to gradient descent. This discovery opened up a new realm of possibilities in analyzing the convergence of Walrasian equilibrium processes. For example, it allowed Cheung et al. (2018) to generalize the convergence results of proportional response dynamics to Fisher markets for buyers with mixed CES utilities. This same idea was applied by Cheung et al. (2013) to prove the convergence of *tâtonnement* in Leontief Fisher markets, using the equivalence between mirror descent (Boyd et al., 2004) on the dual of the Eisenberg-Gale program and *tâtonnement*, first observed by Devanur et al. (2008). More recently, Gao and Kroer (2020) developed methods to solve the Eisenberg-Gale convex program in the case of linear, quasi-linear, and Leontief Fisher markets.

An alternative to the (global) competitive economy model, in which an agent's trading partners are unconstrained, is the Kakade et al. (2004) model of a graphical economies. This model features local markets, in which each agent can set its own prices for purchase only by neighboring agents, and likewise can purchase only from neighboring agents. Auction-like price-adjustment processes have been shown to converge in variants of this model assuming WGS (Andrade et al., 2021).

5.4 Price Adjustment Processes for Walrasian Equilibrium

The most common class of algorithms for computing a Walrasian equilibrium are first-order price adjustment processes simply called **price adjustment processes** (Papadimitriou and Yannakakis, 2010).

Definition 5.4.1 [Price-adjustment process]**.**

Given a Walrasian economy $(m, \boldsymbol{z})$ and an initial price vector $\boldsymbol{p}^{(0)} \in \mathbb{R}_+^m$, a **price adjustment process** $\boldsymbol{\pi}$ consists of an update function $\boldsymbol{\pi} : \bigcup_{\tau \geq 1} (\mathbb{R}_+^m \times \mathbb{R}^m) \to \mathbb{R}_+^m$ that generates a sequence of prices $\{\boldsymbol{p}^{(t)}\}_t$ given by: for all $t = 0, 1, \ldots$,

$$\boldsymbol{p}^{(t+1)} \doteq \boldsymbol{\pi}\left(\bigcup_{k=0}^{t} (\boldsymbol{p}^{(k)}, \boldsymbol{z}(\boldsymbol{p}^{(k)}))\right)$$

An important class of price-adjustment processes are *natural* price-adjustment processes. Intuitively, these are price-adjustment processes where the price of each commodity is updated using only information about the past prices of the commodity itself and its excess demand. These processes are natural in the sense that the price of each commodity is updated only with information relevant to it, and as such if each commodity is sold by a fictional seller, the seller can update the price of its good without having to coordinate with other sellers.

Definition 5.4.2 [Natural Price-Adjustment Process]**.**

Given a Walrasian economy $(m, \boldsymbol{z})$ and an initial price vector $\boldsymbol{p}^{(0)} \in \mathbb{R}_+^m$, a price adjustment process $\boldsymbol{\pi}$ is said to be **natural** if for all commodities, the price adjustment process can be written as as $\boldsymbol{\pi} \doteq (\pi_1, \ldots, \pi_m)$, where for all commodities $j \in [m]$, $\pi_j : \bigcup_{\tau \geq 1} (\mathbb{R}_+ \times \mathbb{R}) \to \mathbb{R}_+^m$ s.t. for all $t = 0, 1, \ldots$,

$$p_j^{(t+1)} \doteq \pi_j\left(\bigcup_{k=0}^{t} (p_j^{(k)}, z_j(\boldsymbol{p}^{(k)}))\right)$$

The canonical type of natural price adjustment processes are *tâtonnement processes* (Walras, 1896; Arrow and Hurwicz, 1958).

Definition 5.4.3 [*Tâtonnement* process]**.**

A *tâtonnement process* is a natural price adjustment process $\boldsymbol{\pi} \doteq (\pi_1, \ldots, \pi_m)$ s.t. for all

$j \in [m]$ and $t \in \mathbb{N}_{++}$, there exists a function $g : \mathbb{R}_+ \times \mathbb{R} \to \mathbb{R}$ that satisfies:

$$\pi_j \left(\bigcup_{k=0}^{t} (p_j^{(k)}, z_j(\boldsymbol{p}^{(k)})) \right) \doteq g(p_j^{(t)}, z_j(\boldsymbol{p}^{(t)})) \tag{5.7}$$

Remark 5.4.1 [On *tâtonnement*]**.**

The verb *tâtonner* is a French word that means to search by trial and error, often connoting a sense of blindness, as the search relies solely on local information. Accordingly, the noun form *tâtonnement* describes a heuristic search process based on trial and error. The *tâtonnement* process is a memoryless price adjustment mechanism where each commodity's next price is determined solely by its current price and excess demand. The term *tâtonnement* is thus aptly chosen, as the price search for each commodity is heuristic, ignoring both past prices and excess demands as well as the current prices and excess demands of other goods.

Traditionally, g is further restricted to be sign preserving, i.e., $\forall p \in \mathbb{R}_+, z \in \mathbb{R}, \mathrm{sign}(g(p, z)) = \mathrm{sign}(z)$, as with this restriction in place a *tâtonnement* process can be seen a mathematical model of the law supply and demand, which stipulates that the price of any commodity in the economy that is demanded (respectively, supplied) in excess will rise (respectively, decrease) (Walras, 1896; Arrow and Hurwicz, 1958).

Now notice that the mirror gradient method applied to the VI $(\mathbb{R}^m_+, -\mathcal{Z})$ defines a family of *tâtonnement* processes parametrized by the kernel function h, which we will call the mirror *tâtonnement* process. While continuous-time variants of the *tâtonnement processes* are known to converge in Arrow-Debreu economies for which the excess demand $\boldsymbol{z}$ is for instance monotone (in which case the excess demand is said to satisfy the **law of demand**, see Definition 5.4.8), as Example 4.3.1 shows the mirror *tâtonnement* process is not guaranteed to converge in such economies.[4] Nevertheless, recall that we can instead apply the mirror extragradient algorithm to the VI $(\mathbb{R}^m_+, -\mathcal{Z})$, which as we have shown, can be guaranteed to converge in VIs that satisfy the Minty condition using the tools developed in Chapter 4. To this end, we introduce the class of variationally stable Walrasian economies.

[4]While this example is presented for VIs, by the equivalence between VIs and Walrasian economies, it also applies to Walrasian economies.

Definition 5.4.4 [Variationally Stable Walrasian Economies]**.**

A Walrasian economy $(m, \mathcal{Z})$, is said to be **variationally stable** on $\mathcal{P} \subseteq \mathbb{R}_+^m$ iff there exists $\boldsymbol{p}^* \in \mathcal{P}$ s.t. for all $\boldsymbol{p} \in \mathcal{P}$ and $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$,

$$\langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p}^* - \boldsymbol{p} \rangle \geq 0$$

If a balanced economy is variationally stable on Δ_m, then we refer to the economy simply as **variationally stable**.

To understand the variational stability condition, consider a fictional auctioneer who buys the commodities sold in the economy and sells them back at prices $\boldsymbol{p} \in \mathbb{R}_+^m$. The profit of the auctioneer for her transaction is given by $\langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p} \rangle$. Now suppose that the auctioneer were to change the prices at which she bought and sold her commodities to prices $\boldsymbol{p}^* \in \mathbb{R}_+^m$, while fixing the quantities of goods bought and sold at the excess demand $\boldsymbol{z}(\boldsymbol{p})$. Then the auctioneer's profit would be given by $\langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p}^* \rangle - \langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p} \rangle = \langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p}^* - \boldsymbol{p} \rangle$. The Minty condition requires the existence of a price vector $\boldsymbol{p}^* \in \mathbb{R}_+^m$, which in hindsight looks to the auctioneer like a more profitable price vector than $\boldsymbol{p}$.

5.4.1 Computation of Walrasian Equilibrium in Balanced Economies

A surprising and important result, which is described by the following (new) lemma, is that the VI $([0,1]^m, -\mathcal{Z})$ associated with any balanced economy $(m, \mathcal{Z})$ satisfies the Minty condition, i.e., any balanced economy is variationally stable on $[0,1]^m$.

Lemma 5.4.1 [Balanced Economies are Variationally Stable on the Unit Box]**.**

Any balanced economy $(m, \mathcal{Z})$ is variationally stable on $[0,1]^m$. In particular, letting $\boldsymbol{p}^* \doteq \boldsymbol{0}_m$, for all prices $\boldsymbol{p} \in [0,1]^m$ and $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$, we have $\langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p}^* - \boldsymbol{p} \rangle \geq 0$.

Proof of Lemma 5.4.1

Let $(m, \mathcal{Z})$ be a balanced economy. Setting $\boldsymbol{p}^* \doteq \boldsymbol{0}_m$, we have:

$$
\begin{aligned}
\langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p}^* - \boldsymbol{p} \rangle &= \langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{0}_m - \boldsymbol{p} \rangle \\
&= \underbrace{\langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{0}_m \rangle}_{=0} - \langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p} \rangle \\
&= -\underbrace{\langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p} \rangle}_{\leq 0} \\
&\geq 0 \enspace ,
\end{aligned}
$$

where the last line follow from weak Walras' law, which is assumed to hold in balanced economies.

This lemma is highly surprising, as it suggests that in balanced economies, which include among others Arrow-Debreu competitive economies (see Chapter 10 for additional details), under suitable continuity assumptions, first-order methods for the VI $([0, 1]^m, -\mathcal{Z})$ are guaranteed to converge to a strong solution.

Hence, with Lemma 5.4.1 in hand, we now turn our attention to solving the VI $([0, 1]^m, -\mathcal{Z})$—or rather, the VI $([0, 1]^m, -\boldsymbol{z})$, since we assume for our algorithms that the excess demand is singleton-valued—and hence computing a Walrasian equilibrium with the mirror extragradient method. Solving the VI $([0, 1]^m, -\boldsymbol{z})$ with the mirror extragradient method, gives rise to a family of price adjustment processes parameterized by the kernel function h which we will call the **mirror extratâtonnement process**.

Remark 5.4.2 [Mirror *Extratâtonnement* is a Natural Price-Adjustment Process]**.**
For the choice of a price space $\mathcal{P} \doteq [0, 1]^m$, and any choice of kernel function s.t. $h(\boldsymbol{p}) \doteq \sum_{j[m]} h_j(p_j)$ for some $\{h_j : \mathbb{R}^m \to \mathbb{R}\}_{j \in [m]}$, the mirror *extratâtonnement* updates can be

Algorithm 6 Mirror Extratâtonnement Process

Input: $m, \boldsymbol{z}, \tau, \eta, h, \mathcal{P}, \boldsymbol{p}^{(0)}$

Output: $\{\boldsymbol{p}^{(t)}\}_{t\in[\tau]}$

1: **for** $t = 1, \ldots, \tau$ **do**

2: $\quad \boldsymbol{p}^{(t+0.5)} \leftarrow \arg\min_{\boldsymbol{p}\in\mathcal{P}} \left\{ \left\langle \boldsymbol{z}(\boldsymbol{p}^{(t)}), \boldsymbol{p}^{(t)} - \boldsymbol{p} \right\rangle + \frac{1}{2\eta}\mathrm{div}_h(\boldsymbol{p}, \boldsymbol{p}^{(t)}) \right\}$

3: $\quad \boldsymbol{p}^{(t+1)} \leftarrow \arg\min_{\boldsymbol{p}\in\mathcal{P}} \left\{ \left\langle \boldsymbol{z}(\boldsymbol{p}^{(t+0.5)}), \boldsymbol{p}^{(t)} - \boldsymbol{p} \right\rangle + \frac{1}{2\eta}\mathrm{div}_h(\boldsymbol{p}, \boldsymbol{p}^{(t)}) \right\}$

return $\{\boldsymbol{p}^{(t+0.5)}\}_{t\in[\tau]}$

written for all commodities $j \in [m]$ and $t \in \mathbb{N}$ as:

$$\begin{aligned} p_j^{(t+0.5)} &\leftarrow \arg\min_{p_j\in[0,1]} \left\{ \left\langle z_j(\boldsymbol{p}^{(t)}), p_j^{(t)} - p_j \right\rangle + \frac{1}{2\eta}\mathrm{div}_{h_j}(p_j, p_j^{(t)}) \right\} \\ p_j^{(t+1)} &\leftarrow \arg\min_{p_j\in[0,1]} \left\{ \left\langle z_j(\boldsymbol{p}^{(t+0.5)}), p_j^{(t)} - p_j \right\rangle + \frac{1}{2\eta}\mathrm{div}_{h_j}(p_j, p_j^{(t)}) \right\} \quad . \end{aligned}$$

Further, multiplying the index of the sequence of price iterates, we can rewrite the above update rule for all commodities $j \in [m]$ and $t \in \mathbb{N}$ as:

$$\begin{aligned} p_j^{(t+1)} &\leftarrow \arg\min_{p_j\in[0,1]} \left\{ \left\langle z_j(\boldsymbol{p}^{(t)}), p_j^{(t)} - p_j \right\rangle + \frac{1}{2\eta}\mathrm{div}_{h_j}(p_j, p_j^{(1)}) \right\} \\ p_j^{(t+2)} &\leftarrow \arg\min_{p_j\in[0,1]} \left\{ \left\langle z_j(\boldsymbol{p}^{(t+1)}), p_j^{(t)} - p_j \right\rangle + \frac{1}{2\eta}\mathrm{div}_{h_j}(p_j, p_j^{(t)}) \right\} \quad . \end{aligned}$$

That is, the mirror *extratâtonnement* process applies a *tâtonnement* update to the current time-step's prices, while on even time-steps, it applies a *tâtonnement* update to the *previous* time-step's prices. As such, the mirror *extratâtonnement* process can be interpreted as a natural price adjustment process.

With the mirror *extrâtonnement* process, and Lemma 5.4.1 in hand, we can apply Theorem 4.3.1 to prove the polynomial-time convergence of the mirror *extrâtonnement* process (Algorithm 6).

Theorem 5.4.1 [Convergence of Mirror *Extratâtonnement*]**.**

Consider the mirror *extrâtonnement* process run on the balanced economy $(m, \boldsymbol{z})$ with a 1-strongly-convex and κ-Lipschitz-smooth kernel function h, any time horizon $t \dot{\in} \mathbb{N}$, any step size $\eta > 0$, a price space $\mathcal{P} \doteq [0,1]^m$, and any initial price vector $\boldsymbol{p}^{(0)} \in [0,1]^m$. Let $\{\boldsymbol{p}^{(k)}, \boldsymbol{p}^{(k+0.5)}\}_t$ be the sequence of prices generated, and suppose there exists $\lambda \in (0, \frac{1}{\sqrt{2}\eta}]$ s.t. $\frac{1}{2}\left\|\boldsymbol{z}(\boldsymbol{p}^{(k+0.5)}) - \boldsymbol{z}(\boldsymbol{p}^{(k)})\right\|^2 \leq \lambda^2 \mathrm{div}_h(\boldsymbol{p}^{(k+0.5)}, \boldsymbol{p}^{(k)})$. If $\boldsymbol{p}^{(\tau)}_{\text{best}} \in \arg\min_{\boldsymbol{x}^{(k+0.5)}:k=0,\ldots,\tau} \mathrm{div}_h(\boldsymbol{p}^{(k+0.5)}, \boldsymbol{p}^{(k)})$, then for some time horizon $\tau \dot{\in} O(\frac{\kappa^2 m^2 \mathrm{div}_h(\boldsymbol{0}_m, \boldsymbol{p}^{(0)})}{\eta^2 \varepsilon^2})$, $\boldsymbol{p}^{(\tau)}_{\text{best}}$ is an ε-Walrasian equilibrium of $(m, \boldsymbol{z})$. Furthermore, $\lim_{t\to\infty} \boldsymbol{p}^{(t+0.5)} = \lim_{t\to\infty} \boldsymbol{p}^{(t)} = \boldsymbol{p}^*$ is a Walrasian equilibrium of $(m, \boldsymbol{z})$.

Proof of Theorem 5.4.1

Since $(m, \boldsymbol{z})$ is a balanced economy, by Lemma 5.4.1, $(m, \boldsymbol{z})$ is variationally stable on $[0,1]^m$, and hence the VI $([0,1]^m, -\mathcal{Z})$ satisfies the Minty condition. Hence, as the mirror *extratâtonnement* process is simply the mirror extragradient method run on the VI $([0,1]^m, -\mathcal{Z})$, the assumptions of Theorem 4.3.1 are satisfied, and we obtain the result.

The convergence guarantee provided by Theorem 5.4.1 is highly general, and does not require Lipschitz-continuity of the excess demand $\boldsymbol{z}$. Rather, the above theorem requires a notion "Bregman-continuity over trajectories" of the *extratâtonnement* process. This broad statement is purposeful, as it is in general not possible to guarantee Lipschitz-continuity of the excess demand in balanced economies. Indeed, the only balanced economies with a Lipschitz-continuous excess demand function are those with a constant excess demand function.

To see this, suppose that $\boldsymbol{z}$ is λ-Lipschitz-continuous on $[0,1]^m$. By homogeneity of degree 0, we have, for all $\alpha > 0$ and $\boldsymbol{p}, \boldsymbol{q} \in [0,1]^m$, $\|\boldsymbol{z}(\boldsymbol{p}) - \boldsymbol{z}(\boldsymbol{q})\| = \|\boldsymbol{z}(\alpha\boldsymbol{p}) - \boldsymbol{z}(\alpha\boldsymbol{q})\| \leq \lambda\alpha \|\boldsymbol{p} - \boldsymbol{q}\|$. Hence, taking $\alpha \to 0$, we have, for all $\boldsymbol{p}, \boldsymbol{q} \in [0,1]^m$, $\boldsymbol{z}(\boldsymbol{q}) = \boldsymbol{z}(\boldsymbol{p})$.

Nevertheless, while Lipschitz-continuity over $[0,1]^m$ is too restrictive, Lipschitz continuity over paths of the mirror *extratâtonnement* process, which we call **pathwise Lipschitz-**

continuity (i.e., $\left\|\boldsymbol{z}(\boldsymbol{p}^{(k+0.5)}) - \boldsymbol{z}(\boldsymbol{p}^{(k)})\right\| \leq \lambda \left\|\boldsymbol{p}^{(k+0.5)} - \boldsymbol{p}^{(k)}]\right\|$)[5] seems to be a mild assumption that holds in a large class of Walrasian economies, as suggested by the experiments reported in Section 5.4.3.

For choices of kernel functions h s.t. the associated Bregman divergence div_h is *not* homogeneous of degree $\alpha > 0$ (i.e., for all $\boldsymbol{p}, \boldsymbol{q} \in \mathbb{R}_+^m$ and $\alpha, \lambda > 0$, $\mathrm{div}_h(\lambda \boldsymbol{p}, \lambda \boldsymbol{q}) \neq \lambda^\alpha \mathrm{div}_h(\boldsymbol{p}, \boldsymbol{q})$), we define the following novel class, which seems likely to capture a broad swath of Walrasian economies:

Definition 5.4.5 [Bregman-continuous Economies]**.**
Given a modulus of continuity $\lambda \geq 0$ and a kernel function $h : \mathcal{P} \to \mathbb{R}$, a (λ, h)**-Bregman-continuous economy** on $\mathcal{P} \subseteq \mathbb{R}_+^m$ is a Walrasian economy $(m, \mathcal{Z})$ whose excess demand correspondence is singleton-valued (i.e., $\mathcal{Z}(\boldsymbol{p}) \doteq \{\boldsymbol{z}(\boldsymbol{p})\}$) and (λ, h)**-Bregman-continuous** on $\mathcal{P}$, i.e., for all $\boldsymbol{p}, \boldsymbol{q} \in \mathcal{P}$,

$$\frac{1}{2} \left\|\boldsymbol{z}(\boldsymbol{p}) - \boldsymbol{z}(\boldsymbol{q})\right\|^2 \leq \lambda^2 \mathrm{div}_h(\boldsymbol{q}, \boldsymbol{p})$$

When h is clear from context, we simply say that the economy and the excess demand $\boldsymbol{z}$ are both λ-Bregman continuous on $\mathcal{P}$.

Bregman continuous functions were introduced into the optimization literature in recent years, and have been shown to contain a large number of important function classes which are not continuous (see, for instance, Lu (2019)). Note that when the kernel function h is chosen to be $h(\boldsymbol{p}) \doteq \frac{1}{2}\|\boldsymbol{p}\|^2$, λ-Bregman-continuity reduces to λ-Lipschitz continuity. Furthermore, the literature on algorithmic general equilibrium theory has considered variants of Bregman continuity to prove the polynomial-time convergence of algorithms to Walrasian equilibria (see, for instance Cheung et al. (2013) and Cheung et al. (2018)). As such, Bregman continuity seems a natural assumption to adopt here.

[5] Notice that when the kernel function in the statement of Theorem 5.4.1 is chosen to be $h(\boldsymbol{p}) \doteq \frac{1}{2}\left\|\boldsymbol{p}\right\|^2$, the condition $\frac{1}{2}\left\|\boldsymbol{z}(\boldsymbol{p}^{(k+0.5)}) - \boldsymbol{z}(\boldsymbol{p}^{(k)})\right\|^2 \leq \lambda^2 \mathrm{div}_h(\boldsymbol{p}^{(k+0.5)}, \boldsymbol{p}^{(k)})$ reduces to $\|\boldsymbol{z}(\boldsymbol{p}^{(k+0.5)}) - \boldsymbol{z}(\boldsymbol{p}^{(k)})\| \leq \lambda \left\|\boldsymbol{p}^{(k+0.5)} - \boldsymbol{p}^{(k)}]\right\|$.

With the definition of Bregman continuous economies in hand, we obtain the following corollary of Theorem 5.4.1.

Corollary 5.4.1 [Convergence of Mirror *Extrâtonnement* under Bregman Continuity]**.**
For some modulus of continuity $\lambda > 0$ and kernel function h, let $(m, \boldsymbol{z})$ be a balanced economy that is (λ, h)-Bregman-continuous on $[0,1]^m$. Consider the mirror *extrâtonnement* process run on $(m, \boldsymbol{z})$, with a 1-strongly-convex and κ-Lipschitz-smooth kernel function h, any time horizon $t \doteq \mathbb{N}$, any step size $\eta \in (0, \frac{1}{\sqrt{2}\lambda}]$, a price space $\mathcal{P} \doteq [0,1]^m$, and any initial price vector $\boldsymbol{p}^{(0)} \in [0,1]^m$. The outputs $\{\boldsymbol{p}^{(t)}, \boldsymbol{p}^{(t+0.5)}\}_t$ satisfy the following: If $\boldsymbol{p}^{(\tau)}_{\text{best}} \in \arg\min_{\boldsymbol{x}^{(k+0.5)}:k=0,\ldots,\tau} \text{div}_h(\boldsymbol{p}^{(k+0.5)}, \boldsymbol{p}^{(k)})$, then for some time horizon $\tau \doteq O(\frac{\kappa^2 m^2 \text{div}_h(\boldsymbol{0}_m, \boldsymbol{p}^{(0)})}{\eta^2 \varepsilon^2})$, $\boldsymbol{p}^{(\tau)}_{\text{best}}$ is an ε-Walrasian equilibrium. Furthermore, $\lim_{t\to\infty} \boldsymbol{p}^{(t+0.5)} = \lim_{t\to\infty} \boldsymbol{p}^{(t)} = \boldsymbol{p}^*$ is a Walrasian equilibrium of $(m, \boldsymbol{z})$.

While these convergence results are useful, it is not immediately clear what types of excess demand functions satisfy Bregman-continuity. As a result, to characterize the Bregman-continuity properties of Walrasian economies, we introduce the following economic parameters, which have been used extensively in the analysis of algorithms to compute Walrasian equilibrium (see, for instance, Cole and Fleischer (2008)).

Definition 5.4.6 [Function Elasticity]**.**
Given any function $\boldsymbol{f} : \mathbb{R}^n \to \mathbb{R}^m$, we define the **elasticity** $\epsilon_{f_j, p_k} : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ of output f_j w.r.t. input x_k between any two points $\boldsymbol{x} \in \mathbb{R}^n$ and $\boldsymbol{y} \in \mathbb{R}^n$ as the percentage change in f_j that results from a one percent change from x_k to y_k:

$$\epsilon_{f_j, x_k}(\boldsymbol{x}, \boldsymbol{y}) \doteq \frac{f_j(\boldsymbol{y}) - f_j(\boldsymbol{x})}{f_j(\boldsymbol{x})} \frac{x_k}{y_k - x_k} \tag{5.8}$$

Overloading notation, we also define the **instantaneous elasticity** as follows:

$$\epsilon_{f_j, x_k}(\boldsymbol{x}, \boldsymbol{y}) \doteq \lim_{h\to 0} \frac{\epsilon_{f_j, x_k}(\boldsymbol{x}, \boldsymbol{x} + h\boldsymbol{j}_k)}{h} = \frac{\partial_{x_k} f_j(\boldsymbol{x}) x_k}{f_j(\boldsymbol{x})} \tag{5.9}$$

Definition 5.4.7 [Elastic Economies]**.**
Given $\overline{\epsilon} \geq 0$, an $\overline{\epsilon}$**-elastic economy** $(m, \boldsymbol{d}, \boldsymbol{s})$ is a Walrasian economy $(m, \mathcal{Z})$ with an **aggregate demand function** $\boldsymbol{d} : \mathbb{R}^m_+ \to \mathbb{R}^m_+$ and **aggregate supply function** $\boldsymbol{s} : \mathbb{R}^m_+ \to \mathbb{R}^m_+$ s.t.

$\mathcal{Z}(\boldsymbol{p}) \doteq \{\boldsymbol{d}(\boldsymbol{p}) - \boldsymbol{s}(\boldsymbol{p})\}$, and for which the following two bounds on the elasticity of the supply and demand hold:

$$\max_{\substack{\boldsymbol{p},\boldsymbol{q}\in\mathbb{R}_+^m \\ j,k\in[m]}} \left|\epsilon_{d_j,p_k}(\boldsymbol{p},\boldsymbol{q})\right| \leq \bar{\epsilon}, \qquad \max_{\substack{\boldsymbol{p},\boldsymbol{q}\in\mathbb{R}_+^m \\ j,k\in[m]}} \left|\epsilon_{s_j,p_k}(\boldsymbol{p},\boldsymbol{q})\right| \leq \bar{\epsilon}$$

The following lemma demonstrates that the excess demand of any $\bar{\epsilon}$-economy with bounded aggregate demand and supply is Bregman-continuous w.r.t. to the log-barrier kernel $h(\boldsymbol{p}) = -\sum_{j\in[m]} \log(p_j)$.

Lemma 5.4.2 [Bregman Continuity for Elastic Economies]**.**

If $(m, \boldsymbol{d}, \boldsymbol{s})$ is an $\bar{\epsilon}$-elastic economy, then, for any 1-strongly-convex kernel function $h : \mathbb{R}_+^m \to \mathbb{R}$, the following bound holds:

$$\frac{1}{2}\left\|\boldsymbol{z}(\boldsymbol{q}) - \boldsymbol{z}(\boldsymbol{p})\right\|^2 \leq \left(\frac{\epsilon\left(\|\boldsymbol{d}(\boldsymbol{p})\| + \|\boldsymbol{s}(\boldsymbol{p})\|\right)}{\|\boldsymbol{p}\|_\infty}\right)^2 \mathrm{div}_h(\boldsymbol{q},\boldsymbol{p})$$

Proof of Lemma 5.4.2

By the $\bar{\epsilon}$-elasticity assumption, we have, for all $\boldsymbol{p},\boldsymbol{q} \in \Delta_m$ and $j,k \in [m]$,

$$\begin{aligned}
\left|\frac{d_j(\boldsymbol{q}) - d_j(\boldsymbol{p})}{d_j(\boldsymbol{p})}\frac{p_k}{q_k - p_k}\right| &\leq \epsilon \\
\frac{|d_j(\boldsymbol{q}) - d_j(\boldsymbol{p})|}{|d_j(\boldsymbol{p})|}\frac{|p_k|}{|q_k - p_k|} &\leq \epsilon \\
|d_j(\boldsymbol{q}) - d_j(\boldsymbol{p})| &\leq \frac{\epsilon|d_j(\boldsymbol{p})|}{|p_k|}|q_k - p_k| \\
|d_j(\boldsymbol{q}) - d_j(\boldsymbol{p})|^2 &\leq \frac{\epsilon^2|d_j(\boldsymbol{p})|^2}{(p_k)^2}\,|q_k - p_k|^2
\end{aligned}$$

Summing up over $j \in [m]$, we have, for all $k \in [m]$,

$$\begin{aligned}
\|\boldsymbol{d}(\boldsymbol{q}) - \boldsymbol{d}(\boldsymbol{p})\|^2 &\leq \frac{\epsilon^2\|\boldsymbol{d}(\boldsymbol{p})\|^2}{(p_k)^2}|q_k - p_k|^2 \\
&\leq \frac{\epsilon^2\|\boldsymbol{d}(\boldsymbol{p})\|^2}{(p_k)^2}\,\|\boldsymbol{q} - \boldsymbol{p}\|^2
\end{aligned}$$

Since h is 1-strongly-convex, for all $\boldsymbol{x},\boldsymbol{y} \in \mathcal{X}$, $\mathrm{div}_h(\boldsymbol{x},\boldsymbol{y}) \geq 1/2\|\boldsymbol{x} - \boldsymbol{y}\|^2$. Hence, we have:

$$\|\boldsymbol{d}(\boldsymbol{q}) - \boldsymbol{d}(\boldsymbol{p})\|^2 \leq \frac{2\epsilon^2\|\boldsymbol{d}(\boldsymbol{p})\|^2}{(p_k)^2}\mathrm{div}_h(\boldsymbol{q},\boldsymbol{p})$$

Taking the square root of both sides and then taking a minimum over $k \in [m]$, we have:

$$\begin{aligned}
\|\boldsymbol{d}(\boldsymbol{q}) - \boldsymbol{d}(\boldsymbol{p})\| &\leq \min_{k\in[m]} \frac{\epsilon\|\boldsymbol{d}(\boldsymbol{p})\|}{p_k}\sqrt{2\mathrm{div}_h(\boldsymbol{q},\boldsymbol{p})} \\
&= \frac{\epsilon\|\boldsymbol{d}(\boldsymbol{p})\|}{\max_{k\in[m]} p_k}\sqrt{2\mathrm{div}_h(\boldsymbol{q},\boldsymbol{p})} \\
&= \frac{\epsilon\|\boldsymbol{d}(\boldsymbol{p})\|}{\|\boldsymbol{p}\|_\infty}\sqrt{2\mathrm{div}_h(\boldsymbol{q},\boldsymbol{p})}
\end{aligned}$$

By a similar argument, we also have:

$$\|\boldsymbol{s}(\boldsymbol{q}) - \boldsymbol{s}(\boldsymbol{p})\| \leq \frac{\epsilon\|\boldsymbol{s}(\boldsymbol{p})\|}{\|\boldsymbol{p}\|_\infty}\sqrt{2\mathrm{div}_h(\boldsymbol{q},\boldsymbol{p})}$$

Combining the two bounds, we then have:

$$\begin{aligned}
\|\boldsymbol{z}(\boldsymbol{q}) - \boldsymbol{z}(\boldsymbol{p})\| &= \|\boldsymbol{d}(\boldsymbol{q}) - \boldsymbol{s}(\boldsymbol{q}) - \boldsymbol{d}(\boldsymbol{p}) + \boldsymbol{s}(\boldsymbol{p})\| \\
&\leq \|\boldsymbol{d}(\boldsymbol{q}) - \boldsymbol{d}(\boldsymbol{p})\| + \|\boldsymbol{s}(\boldsymbol{q}) - \boldsymbol{s}(\boldsymbol{p})\| \\
&\leq \frac{\epsilon\|\boldsymbol{d}(\boldsymbol{p})\|}{\|\boldsymbol{p}\|_\infty}\sqrt{2\mathrm{div}_h(\boldsymbol{q},\boldsymbol{p})} + \frac{\epsilon\|\boldsymbol{s}(\boldsymbol{p})\|}{\|\boldsymbol{p}\|_\infty}\sqrt{2\mathrm{div}_h(\boldsymbol{q},\boldsymbol{p})} \\
&\leq \frac{\epsilon\|\boldsymbol{d}(\boldsymbol{p})\| + \|\boldsymbol{s}(\boldsymbol{p})\|}{\|\boldsymbol{p}\|_\infty}\sqrt{2\mathrm{div}_h(\boldsymbol{q},\boldsymbol{p})}
\end{aligned}$$

Squaring both sides and reorganizing yields:

$$\frac{1}{2}\|\boldsymbol{z}(\boldsymbol{q}) - \boldsymbol{z}(\boldsymbol{p})\|^2 \leq \left(\frac{\epsilon\left(\|\boldsymbol{d}(\boldsymbol{p})\| + \|\boldsymbol{s}(\boldsymbol{p})\|\right)}{\|\boldsymbol{p}\|_\infty}\right)^2 \mathrm{div}_h(\boldsymbol{q},\boldsymbol{p})$$

Lemma 5.4.2 suggests that boundedness of the excess demand and a lower bound on prices is sufficient to ensure Bregman-continuity of the excess demand. Boundedness of the excess demand can be ensured in large class of Walrasian economies including Arrow-Debreu economies. While it is not possible to ensure that prices are bounded from below, since Theorem 5.4.1 requires Bregman continuity over the paths of mirror *extratâtonnement*, we have the following corollary of Theorem 5.4.1.

Corollary 5.4.2 [Convergence of Mirror *Extratâtonnement*]**.**

Let $(m, \boldsymbol{d}, \boldsymbol{s})$ be a balanced and $\bar{\epsilon}$-elastic economy. Consider the mirror *extrâtonnement* process run on $(m, \boldsymbol{z})$, with a 1-strongly-convex and κ-Lipschitz-smooth kernel function h, any time horizon $t \dot{\in} \mathbb{N}$, any step size $\eta > 0$, a price space $\mathcal{P} \doteq [0,1]^m$, and any initial price vector $\boldsymbol{p}^{(0)} \in [0,1]^m$. The output sequence $\{\boldsymbol{p}^{(t)}, \boldsymbol{p}^{(t+0.5)}\}_t$ satisfies the following: If $\boldsymbol{p}^{(\tau)}_{\text{best}} \in \arg\min_{\boldsymbol{x}^{(k+0.5)}:k=0,\dots,\tau} \operatorname{div}_h(\boldsymbol{p}^{(k+0.5)}, \boldsymbol{p}^{(k)})$, and if the step size satisfies $\eta \leq \min_{k\in[\tau]} \left\{ \frac{\|\boldsymbol{p}^{(t)}\|_\infty}{\epsilon(\|\boldsymbol{d}(\boldsymbol{p}^{(t)})\| + \|\boldsymbol{s}(\boldsymbol{p}^{(t)})\|)} \right\}$, then for some $\tau \dot{\in} O(\frac{\kappa^2 m^2 \operatorname{div}_h(\boldsymbol{0}_m, \boldsymbol{p}^{(0)})}{\eta^2 \varepsilon^2})$, $\boldsymbol{p}^{(\tau)}_{\text{best}}$ is an ε-Walrasian equilibrium. Furthermore, if the step size instead satisfies $\eta \leq \min_{k\in\mathbb{N}} \left\{ \frac{\|\boldsymbol{p}^{(t)}\|_\infty}{\epsilon(\|\boldsymbol{d}(\boldsymbol{p}^{(t)})\| + \|\boldsymbol{s}(\boldsymbol{p}^{(t)})\|)} \right\}$, then $\lim_{t\to\infty} \boldsymbol{p}^{(t+0.5)} = \lim_{t\to\infty} \boldsymbol{p}^{(t)} = \boldsymbol{p}^*$ is a Walrasian equilibrium.

Beyond Corollary 5.4.2 we are unable to obtain a stronger polynomial-time convergence result for elastic balanced economies. Such a result may be obtainable using a more fine-grained analysis based on a particular kernel function, or perhaps under additional assumptions. We describe some possible directions for future work, and provide an example of how stronger convergence results using a more fine-grained analysis can be obtained using different kernel functions in Chapter 6.

Remark 5.4.3 [Directions for future work]**.**

For balanced economies, one suitable choice of kernel function is the **logistic loss** $h_{\text{LL}}(\boldsymbol{p}) \doteq \sum_{j\in[m]} \left[p_j \log(p_j) + (1-p_j)\log(1-p_j) \right]$ defined on $\mathcal{P} \doteq [0,1]^m$, which defines the **logistic divergence** $\operatorname{div}_{h_{\text{LL}}}(\boldsymbol{p}, \boldsymbol{q}) \doteq \sum_{j\in[m]} \left[p_j \log\left(\frac{p_j}{q_j}\right) + (1-p_j)\log\left(\frac{1-p_j}{1-q_j}\right) \right]$. For this choice of kernel function, the mirror *extratâtonnement* process reduces to the **logistic *extratâtonnement* process**, defined as follows: for all commodities $j \in [m]$ and time horizons $t \in \mathbb{N}$,

$$p_j^{(t+0.5)} \doteq \frac{1}{1 + \frac{1-p_j^{(t)}}{p_j^{(t)}} e^{-\eta z_j(\boldsymbol{p}^{(t)})}}$$
$$p_j^{(t+1)} \doteq \frac{1}{1 + \frac{1-p_j^{(t)}}{p_j^{(t)}} e^{-\eta z_j(\boldsymbol{p}^{(t+0.5)})}}$$

This kernel function has three desirable properties. First, the price updates associated with logistic loss require no projection, and as such the logistic *extratâtonnement* process is a highly natural price adjustment process. Second, logistic divergence is not homogeneous, and as such Bregman continuity w.r.t. logistic loss does not imply that the excess demand is constant. Third, the value of the logistic divergence function tends to infinity as any one of the prices tends to 0, i.e., for all $\boldsymbol{p} \in [0,1]^m$, $\lim_{\boldsymbol{q}\to\boldsymbol{0}} \mathrm{div}_{h_{\mathrm{LL}}}(\boldsymbol{p},\boldsymbol{q}) \to \infty$, which is highly desirable, since for many Walrasian economies, the excess demand when the price of any good heads to 0, is strictly positive.

Remark 5.4.4 [Contributions, and Connection to Impossibility Results]**.**
To the best of our knowledge, Corollary 5.4.2 is the most general convergence guarantee known for (natural) price adjustment processes in Walrasian economies. As we will show in Chapter 10, this result also implies the convergence of price adjustment processes to a Walrasian equilibrium in the canonical class of Walrasian economies known as Arrow-Debreu economies, for which the excess demand is Bregman-continuous. This in turn makes the mirror *extratâtonnement* process the first price adjustment process with a global polynomial-time convergence guarantee to Walrasian equilibrium in Arrow-Debreu economies (i.e., Arrow-Debreu competitive equilibrium), without imposing the highly restrictive weak gross substitutes assumption.

Theorem 5.4.1 might, at first, seem to contradict the impossibility result of Papadimitriou and Yannakakis (2010), which states that for any price adjustment process $\boldsymbol{\pi}$, there exists a balanced economy $(m, \boldsymbol{z})$ that is λ-Lipschitz-continuous on Δ_m for which $\boldsymbol{\pi}$ fails to converge to an ε-Walrasian equilibrium in $\mathrm{poly}(1/\varepsilon)$ evaluations of the excess demand function. However, *we assume that the Walrasian economy is Lipschitz-continuous on* $[0,1]^m$*, rather than* Δ_m. As such, as our continuity assumption is stronger, so our computational result does not contradict the aforementioned impossibility result. Furthermore, our result demonstrates that the computational impossibility results for Walrasian equilibria arise because of "edge case" Walrasian economies in which the excess demand is not sufficiently continuous. We contend that polynomial-time computation of Walrasian equilibrium in balanced economies economies via price adjustment processes is possible.

Our result is also *not* in contradiction with PPAD-hardness results on the computation of Arrow-Debreu equilibrium prices in Leontief Arrow-Debreu economies (Codenotti et al., 2006; Deng and Du, 2008) and additively separable, piecewise linear and concave Arrow-Debreu economies (Chen et al., 2009). Without further assumptions on such economies, the excess demand $\boldsymbol{z}$ can only be shown to be Lipschitz-continuous on Δ_m, and not on $[0,1]^m$.

Beyond the general results obtained in this section, next, to obtain stronger convergence results, we restrict our attention to competitive economies that satisfy WARP.

5.4.2 Computation of Walrasian Equilibrium in Variationally Stable Competitive Economies

In light of Lemma 5.4.2, a sensible question to investigate is whether the mirror *extratâtonnement* process can be guaranteed to converge when the price space is chosen to be $\mathcal{P} \doteq \Delta_m$, thus guaranteeing Bregman-continuity of the excess demand, since $\max_{\boldsymbol{p}\in\Delta} \frac{1}{\max_{k\in[m]} p_k} = \frac{1}{m}$. However, the set of Walrasian equilibria of any balanced economy $(m, \mathcal{Z})$ is not necessarily a subset of the strong solutions of the VI $(\Delta_m, \mathcal{Z})$. Nevertheless, if the economy $(m, \mathcal{Z})$ is assumed to be competitive, then by Theorem 5.2.3, the set of strong solutions of the VI $(\Delta_m, \mathcal{Z})$ *is* a subset of the set of Walrasian equilibria of $(m, \mathcal{Z})$.

As the price space $\mathcal{P} = \Delta_m$ does not include the zero vector $\mathbf{0}_m$, which ensures that balanced economies are variationally stable, the restriction of the price space to Δ_m effectively "destabilizes" the economy, making the computation of a Walrasian equilibrium intractable. To overcome this challenge, we focus our attention on the class of competitive economies that are variationally stable on Δ_m.

Remark 5.4.5 [Interpreting Minty's Condition]**.**
For balanced economies, by weak Walras' law, a sufficient condition for the economy to be variationally stable on Δ_m is the existence of $\boldsymbol{p}^* \in \Delta_m$ s.t. for all prices $\boldsymbol{p} \in \Delta_m$ and $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$,

$$\langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p}^* \rangle \geq 0$$

Now, suppose there exists a commodity $j \in [m]$ that is (weakly) demanded in excess for all $\boldsymbol{p} \in \Delta_m$, i.e., $z_j(\boldsymbol{p}) \geq 0$. Then, setting $\boldsymbol{p}^* = \boldsymbol{j}_j$, we have $\langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p}^* \rangle = \langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{j}_j \rangle = z_j(\boldsymbol{p}) \geq 0$.

Hence, if there is a good that is never supplied in excess, the economy is variationally stable.

Alternatively, a balanced economy is variationally stable on Δ_m whenever there exist two commodities $j, k \in [m]$ whose excess demands are negatively proportional to one another at all prices, i.e., $\exists \alpha > 0$, s.t. $z_j(\boldsymbol{p}) \geq -\alpha z_k(\boldsymbol{p})$. Then, setting $\boldsymbol{p}^* = \frac{1}{1+\alpha}\boldsymbol{j}_j + \frac{\alpha}{(1+\alpha)}\boldsymbol{j}_k$, we have $\langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p}^* \rangle = \frac{1}{1+\alpha} z_j(\boldsymbol{p}) + \frac{\alpha}{(1+\alpha)} z_k(\boldsymbol{p}) \geq \frac{-\alpha}{1+\alpha} z_k(\boldsymbol{p}) + \frac{\alpha}{(1+\alpha)} z_k(\boldsymbol{p}) = 0$. In light of this observation, the variational stability assumption on Δ_m can be seen as a rather mild assumption, as commodities whose excess demands are negatively correlated are abundant in the real world. For instance, electricity and batteries: whenever the excess demand for electricity is positive, this should mean that there is no need to store electricity; that is, the excess demand for batteries is negative.

We now discuss some important classes of Walrasian economies that are variationally stable on Δ_m. The most basic class comprises those economies that satisfy the law of supply and demand. Intuitively, the excess demand in these economies is downward sloping.

Definition 5.4.8 [Law of Supply and Demand Economies]**.**
Given a Walrasian economy $(m, \mathcal{Z})$, an excess demand correspondence is said to satisfy the **law of supply and demand** iff

$$\langle \boldsymbol{z}(\boldsymbol{q}) - \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{q} - \boldsymbol{p} \rangle \leq 0 \qquad \text{for all } \boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p}), \boldsymbol{z}(\boldsymbol{q}) \in \mathcal{Z}(\boldsymbol{q}) \tag{5.10}$$

If the excess demand correspondence of a Walrasian economy satisfies the law of supply and demand, we refer to the economy colloquially as a **law of supply and demand economy**.

We note that the excess demand of a Walrasian economies satisfies the law of supply and demand iff $-\mathcal{Z}$ is monotone. This implies that $-\mathcal{Z}$ is quasimonotone, and hence for any non-empty and compact price space $\mathcal{P} \subseteq \mathbb{R}^m_+$, the VI $(\mathcal{P}, -\mathcal{Z})$ satisfies the Minty condition (see Lemma 3.1 of He (2017)), meaning that any Walrasian economy which satisfies the law of supply and demand is variationally stable on $\mathcal{P}$.

Another important class of Walrasian economies that are variationally stable on Δ_m is the class that satisfies the weak gross substitutes condition. Intuitively, these are Walrasian economies for which the excess demand for a given good can only increase when the price of some other good increases. While we omit the proof as it is involved, we note that any continuous balanced weak gross substitutes Walrasian economy $(m, \mathcal{Z})$ that satisfies Walras' law (i.e., for all $\boldsymbol{p} \in \mathbb{R}_+^m, \boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$ and $\boldsymbol{p} \cdot \boldsymbol{z}(\boldsymbol{p}) = 0$) is a subset of the class of variationally stable economies on $\mathcal{P} \subseteq \mathbb{R}_+^m$, for any non-empty and compact price space $\mathcal{P}$ (see, for instance, Lemma 5 of Arrow et al. (1959)).

Definition 5.4.9 [Weak Gross Substitutes Economies]**.**

Given a Walrasian economy $(m, \mathcal{Z})$, an excess demand correspondence is said to satisfy the **weak gross substitutes condition (WGS)** iff for all $\boldsymbol{p}, \boldsymbol{q} \in \mathbb{R}_+^m$ s.t. for some $k \in [m]$, $q_k > p_k$ and for all $j \neq k, q_j = p_j$, we have:

$$z_j(\boldsymbol{q}) \geq z_j(\boldsymbol{p}) \qquad \text{for all } \boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p}), \boldsymbol{z}(\boldsymbol{q}) \in \mathcal{Z}(\boldsymbol{q}) \tag{5.11}$$

If the above inequality holds strictly, then the excess demand is said to satisfy the gross substitutes condition (GS). Further, if the excess demand correspondence of a Walrasian economy satisfies WGS (respectively, GS), we refer to the economy colloquially as a **WGS (respectively, GS) economy**.

Going further, we can show that any Walrasian economy which satisfies the well-known weak axiom of revealed preferences (Afriat, 1967; Arrow and Hurwicz, 1958) is variationally stable on Δ_m (and, more generally, on any non-empty and compact price space $\mathcal{P} \subseteq \mathbb{R}^m$). To this end, let us first define the weak axiom of revealed preferences for Walrasian economies.

Definition 5.4.10 [WARP excess demand]**.**

Given a Walrasian economy $(m, \mathcal{Z})$, an excess demand correspondence is said to satisfy the **weak axiom of revealed preferences (WARP)** iff for all $\boldsymbol{p}, \boldsymbol{q} \in \mathcal{P}$, $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$, $\boldsymbol{z}(\boldsymbol{q}) \in \mathcal{Z}(\boldsymbol{q})$ with $\boldsymbol{z}(\boldsymbol{p}) \neq \boldsymbol{z}(\boldsymbol{q})$:

$$\langle \boldsymbol{z}(\boldsymbol{q}), \boldsymbol{p} \rangle \leq \langle \boldsymbol{z}(\boldsymbol{q}), \boldsymbol{q} \rangle \text{ implies } \langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{q} \rangle > \langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p} \rangle$$

If the excess demand correspondence of a Walrasian economy satisfies WARP, we refer to the economy colloquially as a **WARP economy**.

Intuitively, a WARP excess demand function is one at which, when the auctioneer is deciding between prices $\boldsymbol{p}$ or prices $\boldsymbol{q}$ so as to maximize profit, if $\boldsymbol{q}$ is weakly more profitable when the excess demand is $\boldsymbol{z}(\boldsymbol{q})$, then $\boldsymbol{q}$ is strictly more profitable when the excess demand is $\boldsymbol{z}(\boldsymbol{p})$.

Remark 5.4.6.

This definition of (WARP) is adapted to arbitrary Walrasian economies, and as such is a generalization of the usual definition for economies that satisfy Walras' law (i.e., for all $\boldsymbol{p} \in \mathbb{R}^m_+$, $\boldsymbol{p} \cdot \boldsymbol{z}(\boldsymbol{p}) = 0$), which requires that $\mathcal{Z}$ is singleton-valued and $\langle \boldsymbol{z}(\boldsymbol{q}), \boldsymbol{p} \rangle \leq 0$ and $\boldsymbol{z}(\boldsymbol{p}) \neq \boldsymbol{z}(\boldsymbol{q}) \implies \langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{q} \rangle > 0$ (i.e., for all $\boldsymbol{p} \in \mathbb{R}^m_+$, $\boldsymbol{p} \cdot \boldsymbol{z}(\boldsymbol{p}) = 0$).

As we show next, WARP implies that $-\mathcal{Z}$ is pseudomonotone in balanced economies.[6]

Lemma 5.4.3 [WARP $\implies$ pseudomonotone]**.**

If the excess demand correspondence $\mathcal{Z}$ of a Walrasian economy $(m, \mathcal{Z})$ satisfies WARP, then its negation $-\mathcal{Z}$, the excess supply is pseudomonotone.

Proof

If $\mathcal{Z}$ satisfies WARP, and $\langle -\boldsymbol{z}(\boldsymbol{q}), \boldsymbol{q} - \boldsymbol{p} \rangle = \langle \boldsymbol{z}(\boldsymbol{q}), \boldsymbol{p} - \boldsymbol{q} \rangle \leq 0$, then $\langle -\boldsymbol{z}(\boldsymbol{p}), \boldsymbol{q} - \boldsymbol{p} \rangle = \langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p} - \boldsymbol{q} \rangle$

If $\boldsymbol{z}(\boldsymbol{p}) \neq \boldsymbol{z}(\boldsymbol{q})$, then, by WARP, $\langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p} - \boldsymbol{q} \rangle < 0$.

Otherwise, if $\boldsymbol{z}(\boldsymbol{p}) = \boldsymbol{z}(\boldsymbol{q})$, then $\langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p} - \boldsymbol{q} \rangle = \langle \boldsymbol{z}(\boldsymbol{q}), \boldsymbol{p} - \boldsymbol{q} \rangle \leq 0$.

That is, if $\mathcal{Z}$ satisfies WARP, then

$$\langle -\boldsymbol{z}(\boldsymbol{q}), \boldsymbol{q} - \boldsymbol{p} \rangle \leq 0 \implies \langle -\boldsymbol{z}(\boldsymbol{p}), \boldsymbol{q} - \boldsymbol{p} \rangle \leq 0$$

Hence, $-\mathcal{Z}$ is pseudomonotone.

[6]To be more precise, we note that an excess demand function $\mathcal{Z}$ satisfies WARP iff $-\mathcal{Z}$ is strictly pseudomonotone. However, as this result will not be used, we present the more general result, which holds for the direction of interest.

An important consequence of Lemma 5.4.3 is that since $-\mathcal{Z}$ is pseudomonotone, for any non-empty and compact price space $\mathcal{P} \subseteq \mathbb{R}^m_+$, the VI $(\mathcal{P}, -\mathcal{Z})$ satisfies the Minty condition (see Lemma 3.1 of He (2017)). As a result, we have the following corollary of Lemma 5.4.3.

Corollary 5.4.3 [WARP $\implies$ Variationally Stable]**.**
Any Walrasian economy that satisfies WARP is variationally stable on any non-empty and compact price space $\mathcal{P} \subseteq \mathbb{R}^m_+$.

To use Lemma 5.4.2, we have to ensure that the excess demand of the economy is bounded, which, as we will show in Chapter 10, is a mild assumption satisfied in all Arrow-Debreu economies—thus necessitating the following definition.

Definition 5.4.11 [Bounded economies]**.**
Given $\overline{z} \geq 0$, a $\overline{z}$**-bounded economy** $(m, \boldsymbol{d}, \boldsymbol{s})$ is a Walrasian economy $(m, \mathcal{Z})$ that consists of an **aggregate demand function** $\boldsymbol{d} : \mathbb{R}^m_+ \to \mathbb{R}^m_+$ and an **aggregate supply function** $\boldsymbol{s} : \mathbb{R}^m_+ \to \mathbb{R}^m_+$ s.t. $\mathcal{Z}(\boldsymbol{p}) \doteq \{\boldsymbol{d}(\boldsymbol{p})\} - \{\boldsymbol{s}(\boldsymbol{p})\}$ and the following bounds hold:

$$\|\boldsymbol{d}\|_\infty \leq \overline{z} \qquad\qquad \|\boldsymbol{s}\|_\infty \leq \overline{z}$$

With these definitions in place, we can now apply Lemma 5.4.2 to prove the polynomial-time convergence of the mirror *extrâtonnement* process in conjunction with Theorem 5.4.1.

Theorem 5.4.2 [Mirror *Extratâtonnement* Convergence in Δ_m]**.**
For any $\overline{\epsilon}, \overline{z} > 0$, let $(m, \boldsymbol{d}, \boldsymbol{e})$ be an $\overline{\epsilon}$-elastic and $\overline{z}$-bounded balanced economy that is variationally stable on Δ_m. Consider the mirror *extrâtonnement* process run on $(m, \boldsymbol{z})$, with a 1-strongly-convex and κ-Lipschitz-smooth kernel function h, any time horizon $\tau \doteq \mathbb{N}$, any step size $\eta \in (0, \frac{1}{2\sqrt{2}m\epsilon\overline{z}}]$, a price space $\mathcal{P} \doteq \Delta_m$, and any initial price vector $\boldsymbol{p}^{(0)} \in \Delta_m$. The output sequence $\{\boldsymbol{p}^{(t)}, \boldsymbol{p}^{(t+0.5)}\}_t$ satisfies the following convergence bound:

$$\min_{k=0,\dots,\tau} \max_{\boldsymbol{p}\in\Delta} \langle \boldsymbol{z}(\boldsymbol{p}^{(k+0.5)}), \boldsymbol{p} - \boldsymbol{p}^{(k+0.5)} \rangle \leq \frac{2\sqrt{2}(1+\kappa)}{\eta} \frac{\sqrt{\max_{\boldsymbol{p}\in\Delta} \mathrm{div}_h(\boldsymbol{p}, \boldsymbol{p}^{(0)})}}{\sqrt{\tau}}$$

Furthermore, $\lim_{t\to\infty} \boldsymbol{p}^{(t+0.5)} = \lim_{t\to\infty} \boldsymbol{p}^{(t)} = \boldsymbol{p}^*$ is a Walrasian equilibrium.

Proof of Theorem 5.4.2

Since $(m, \boldsymbol{z})$ is variationally stable on Δ_m, the VI $(\Delta_m, -\mathcal{Z})$ satisfies the Minty condition. In addition, since by assumption, the economy is $\overline{\epsilon}$-elastic and $\overline{z}$-bounded, by Lemma 5.4.2, $\boldsymbol{z}$ is $(2m\epsilon\overline{z})$-Bregman-continuous on Δ_m. Therefore,

$$
\begin{aligned}
\frac{1}{2}\|\boldsymbol{z}(\boldsymbol{q}) - \boldsymbol{z}(\boldsymbol{p})\|^2 &\leq \left(\frac{\epsilon\left(\|\boldsymbol{d}(\boldsymbol{p})\| + \|\boldsymbol{s}(\boldsymbol{p})\|\right)}{\|\boldsymbol{p}\|_\infty}\right)^2 \operatorname{div}_h(\boldsymbol{q}, \boldsymbol{p}) \\
&\leq \max_{\boldsymbol{p} \in \Delta_m} \left(\frac{\epsilon\left(\|\boldsymbol{d}(\boldsymbol{p})\| + \|\boldsymbol{s}(\boldsymbol{p})\|\right)}{\|\boldsymbol{p}\|_\infty}\right)^2 \operatorname{div}_h(\boldsymbol{q}, \boldsymbol{p}) \\
&\leq \left(\frac{\epsilon\left(\|\boldsymbol{d}\|_\infty + \|\boldsymbol{s}\|_\infty\right)}{\min_{\boldsymbol{p} \in \Delta_m}\|\boldsymbol{p}\|_\infty}\right)^2 \operatorname{div}_h(\boldsymbol{q}, \boldsymbol{p}) \\
&\leq \left(\frac{2\epsilon\overline{z}}{\frac{1}{m}}\right)^2 \operatorname{div}_h(\boldsymbol{q}, \boldsymbol{p}) \\
&\leq (2m\epsilon\overline{z})^2 \operatorname{div}_h(\boldsymbol{q}, \boldsymbol{p}) \quad .
\end{aligned}
$$

That is, the excess demand is Bregman-continuous with a continuity modulus which depends on the number of commodities, the elasticity of the excess demand, and the maximum absolute value of the excess demand.

Now, given the output sequence $\{\boldsymbol{p}^{(t)}, \boldsymbol{p}^{(t+0.5)}\}_t$, let $\boldsymbol{p}^{(\tau)}_{\text{best}} \dot{\in} \arg\min_{\boldsymbol{x}^{(k+0.5)}: k=0,\dots,\tau} \operatorname{div}_h(\boldsymbol{p}^{(k+0.5)}, \boldsymbol{p}^{(k)})$. As mirror *extratâtonnement* is simply the mirror extragradient method run on the VI $(\Delta_m, -\mathcal{Z})$, and as the assumptions of Theorem 4.3.1 are satisfied, we arrive at the following bound:

$$
\min_{k=0,\dots,\tau} \max_{\boldsymbol{p} \in \Delta} \langle -\boldsymbol{z}(\boldsymbol{p}^{(k+0.5)}), \boldsymbol{p}^{(k+0.5)} - \boldsymbol{p} \rangle \leq \frac{2(1+\kappa)\mathrm{diam}(\Delta_m)}{\eta} \frac{\sqrt{\max_{\boldsymbol{p} \in \Delta} \operatorname{div}_h(\boldsymbol{p}, \boldsymbol{p}^{(0)})}}{\sqrt{\tau}}
$$

$$
\min_{k=0,\dots,\tau} \max_{\boldsymbol{p} \in \Delta} \langle \boldsymbol{z}(\boldsymbol{p}^{(k+0.5)}), \boldsymbol{p} - \boldsymbol{p}^{(k+0.5)} \rangle \leq \frac{2\sqrt{2}(1+\kappa)}{\eta} \frac{\sqrt{\max_{\boldsymbol{p} \in \Delta} \operatorname{div}_h(\boldsymbol{p}, \boldsymbol{p}^{(0)})}}{\sqrt{\tau}}
$$

Furthermore, $\lim_t \boldsymbol{p}^{(t)} = \lim_{t \to \infty} \boldsymbol{p}^{t+0.5} = \boldsymbol{p}^*$ is a Walrasian equilibrium.

We make a few remarks before turning our attention to the analysis of the mirror extragradient method in the Scarf economy.

Remark 5.4.7 [Contribution]**.**

In a seminal paper, Arrow and Hurwicz (1958) show that a continuous version of *tâtonnement* converges to a Walrasian equilibrium in Walrasian economies that satisfy WARP. To the best of our knowledge, a corresponding result was not known for discrete-time *tâtonnement*, the variant studied here. As such, ours is the first polynomial-time computation result for ε-Walrasian equilibrium. It is also the first convergence result of any kind for a price adjustment process for the class of Arrow-Debreu economies that satisfy WARP.

Remark 5.4.8 [Boundedness of excess demand]**.**

The assumption that there exists $\overline{z} \geq 0$ s.t. for all $t \in [\tau]$, $\left\|\boldsymbol{z}(\boldsymbol{p}^{(t)})\right\| \leq \overline{z}$ is a common assumption in the analysis of discrete-time price adjustment processes (see, for instance, Cheung et al. (2013) or Chapter 6), and is often guaranteed by a more fine-grained analysis of the Walrasian economy at hand. We present Theorem 5.4.2 in this format to maintain generality for future work.

Remark 5.4.9 [Local convergence of mirror *extratâtonnement*]**.**

The local convergence behavior of mirror *extratâtonnement* can similarly be inferred from Theorem 5.4.2 by instead applying Theorem 4.3.2, and replacing the assumption that the Arrow-Debreu satisfies the Minty condition with the assumption that the initial price iterate starts close enough to a local Minty solution.

Mirror *Extratâtonnement* in Scarf Economies

One of the earliest negative and most discouraging results in the literature on price-adjustment processes is an example of a Walrasian economy provided by Herbert Scarf in which continuous-time *tâtonnement* cycles around the Walrasian equilibrium of the economy, while discrete-time variants spiral away from the equilibrium, starting from any non-equilibrium initial conditions.

Definition 5.4.12.

A **Scarf economy** $\boldsymbol{z}^{\text{scarf}}$ is a Walrasian economy $(3, \boldsymbol{z}^{\text{scarf}})$ with 3 goods for which the excess demand is singleton-valued and given by the function:

$$\boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}) \doteq \begin{pmatrix} \frac{p_1}{p_1+p_2} + \frac{p_3}{p_1+p_3} - 1 \\ \frac{p_1}{p_1+p_2} + \frac{p_2}{p_2+p_3} - 1 \\ \frac{p_2}{p_2+p_3} + \frac{p_3}{p_1+p_3} - 1 \end{pmatrix}$$

The following lemma summarizes properties of the Scarf economy.

Lemma 5.4.4 [Properties of the Scarf Economy]**.**

The Scarf economy is a balanced economy that satisfies Walras' law, i.e., for all $\boldsymbol{p} \in \mathbb{R}_+^m$, $\boldsymbol{p} \cdot \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}) = 0$. Further, the set of Walrasian equilibrium of the Scarf economy $\boldsymbol{z}^{\text{scarf}}$ is given by $\mathcal{WE}(\boldsymbol{z}^{\text{scarf}}) \doteq \{\lambda \mathbf{1}_3 \mid \lambda > 0\}$.

Proof of Lemma 5.4.4

First, notice that the Scarf economy is homogeneous of degree 0. That is, for all $\lambda \geq 0$, we have:

$$\boldsymbol{z}^{\text{scarf}}(\lambda\boldsymbol{p}) \doteq \begin{pmatrix} \frac{\lambda p_1}{\lambda p_1+\lambda p_2} + \frac{\lambda p_3}{\lambda p_1+\lambda p_3} - 1 \\ \frac{\lambda p_1}{\lambda p_1+\lambda p_2} + \frac{\lambda p_2}{\lambda p_2+\lambda p_3} - 1 \\ \frac{\lambda p_2}{\lambda p_2+\lambda p_3} + \frac{\lambda p_3}{\lambda p_1+\lambda p_3} - 1 \end{pmatrix} = \begin{pmatrix} \frac{p_1}{p_1+p_2} + \frac{p_3}{p_1+p_3} - 1 \\ \frac{p_1}{p_1+p_2} + \frac{p_2}{p_2+p_3} - 1 \\ \frac{p_2}{p_2+p_3} + \frac{p_3}{p_1+p_3} - 1 \end{pmatrix} = \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p})$$

Second, for all $\boldsymbol{p} \in \mathbb{R}^m$, we have:

$$\begin{aligned} \boldsymbol{p} \cdot \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}) &= \frac{p_1^2}{p_1+p_2} + \frac{p_1 p_3}{p_1+p_3} - p_1 + \frac{p_1 p_2}{p_1+p_2} + \frac{p_2^2}{p_2+p_3} - p_2 + \frac{p_2 p_3}{p_2+p_3} + \frac{p_3^2}{p_1+p_3} - p_3 \\ &= \frac{p_1^2 + p_1 p_2}{p_1+p_2} + \frac{p_2^2 + p_2 p_3}{p_2+p_3} + \frac{p_3^2 + p_1 p_3}{p_1+p_3} - p_1 - p_2 - p_3 \\ &= \frac{p_1(p_1+p_2)}{p_1+p_2} + \frac{p_2(p_2+p_3)}{p_2+p_3} + \frac{p_3(p_3+p_1)}{p_1+p_3} - p_1 - p_2 - p_3 \\ &= 0 \end{aligned}$$

Finally, observe that for $\boldsymbol{p}^* = \mathbf{1}_m$, we have $\boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}^*) = \mathbf{0}_m$, and thus, $\boldsymbol{p}^* \cdot \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}^*) = 0$. Notice that this equilibrium is unique up to positive scaling since if the price of any commodity is changed from $\boldsymbol{p}^*$, then the excess demand for another commodity is guaranteed to decrease while the excess demand of some other commodity increases.

Our next result shows that the Scarf economy is variationally stable and Lipschitz-continuous for any suitably chosen price space.

Lemma 5.4.5 [Variational Stability and Bregman-continuity of the Scarf Economy]**.** Any Scarf economy $\boldsymbol{z}^{\text{scarf}}$ is variationally stable on Δ_m. Further, for any $\underline{p} \in (0, 1/3)$ and any 1-strongly-convex kernel function $h : \mathbb{R}^3_+ \to \mathbb{R}$, the Scarf economy $\boldsymbol{z}^{\text{scarf}}$ is variationally stable and $(\frac{3}{\underline{p}^2}, h)$-Bregman-continuous on $[\underline{p}, 1]^3$.

Proof of Lemma 5.4.5

Part 1: Variational stability on Δ_m. We claim that for $\boldsymbol{p}^* = (1/3, 1/3, 1/3)$, the Scarf economy is variationally stable on Δ_m, i.e., for all prices $\boldsymbol{p} \in \Delta_3$ and all $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$, we have $\left\langle \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}), \boldsymbol{p}^* - \boldsymbol{p} \right\rangle \geq 0$.

First, notice that expanding the expression $\left\langle \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}), \boldsymbol{p}^* - \boldsymbol{p} \right\rangle$ yields: for all $\boldsymbol{p} \in \Delta_m$,

$$\begin{aligned}
\left\langle \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}), \boldsymbol{p}^* - \boldsymbol{p} \right\rangle &= \left\langle \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}), \boldsymbol{p}^* \right\rangle - \underbrace{\left\langle \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}), \boldsymbol{p} \right\rangle}_{=0} \\
&= \left\langle \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}), \boldsymbol{p}^* \right\rangle \\
&= 2\frac{p_1}{p_1 + p_2} + 2\frac{p_2}{p_2 + p_3} + 2\frac{p_3}{p_1 + p_3} - \underbrace{1/3 - 1/3 - 1/3}_{=-1} \\
&= 2\frac{p_1}{p_1 + p_2} + 2\frac{p_2}{p_2 + p_3} + 2\frac{p_3}{p_1 + p_3} - 1
\end{aligned}$$

We proceed by cases.

Case 1 $p_1 \geq p_2$**.**

$$\begin{aligned}
\left\langle \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}), \boldsymbol{p}^* - \boldsymbol{p} \right\rangle &= 2\frac{p_1}{p_1 + p_2} + \underbrace{2\frac{p_2}{p_2 + p_3}}_{\geq 0} + \underbrace{2\frac{p_3}{p_1 + p_3}}_{\geq 0} - 1 \\
&\geq 2\frac{p_1}{p_1 + \underbrace{p_2}_{\leq p_1}} - 1 \\
&\geq 2\frac{p_1}{p_1 + p_1} - 1 \\
&= 1 - 1 \\
&= 0
\end{aligned}$$

Case 2 $p_1 < p_2$**.**

$$\begin{aligned}\left\langle \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}), \boldsymbol{p}^* - \boldsymbol{p} \right\rangle &= \underbrace{2\frac{p_1}{p_1+p_2}}_{\geq 0} + 2\frac{p_2}{p_2+p_3} + 2\frac{p_3}{p_1+p_3} - 1 \\ &= 2\frac{p_2}{p_2+p_3} + 2\frac{p_3}{\underbrace{p_1}_{<p_2} + p_3} - 1 \\ &= 2\frac{p_2}{p_2+p_3} + 2\frac{p_3}{p_2+p_3} - 1 \\ &= 2\frac{p_2+p_3}{p_2+p_3} - 1 \\ &= 2 - 1 \\ &\geq 0\end{aligned}$$

Hence, it holds that $\left\langle \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}), \boldsymbol{p}^* - \boldsymbol{p} \right\rangle \geq 0$, and the Scarf economy is variationally stable on Δ_m.

Part 2: Variational stability and Bregman-continuity on $[\underline{p}, 1]^3$**.** First, for variational stability on $[\underline{p}, 1]$, observe that the proof for Case 1 applies directly, by simply replacing Δ_m by $[\underline{p}, 1]$.

Second, notice that the Scarf excess demand is differentiable, with its Jacobian matrix given by:

$$\nabla \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p}) = \begin{bmatrix} -\frac{p_1}{(p_1+p_2)^2} - \frac{p_3}{(p_1+p_3)^2} & -\frac{p_1}{(p_1+p_2)^2} & -\frac{p_3}{(p_1+p_3)^2} \\ -\frac{p_1}{(p_1+p_2)^2} & -\frac{p_1}{(p_1+p_2)^2} - \frac{p_3}{(p_2+p_3)^2} & -\frac{p_2}{(p_2+p_3)^2} \\ -\frac{p_3}{(p_1+p_3)^2} & -\frac{p_3}{(p_2+p_3)^2} & -\frac{p_2}{(p_2+p_3)^2} - \frac{p_3}{(p_1+p_3)^2} \end{bmatrix}$$

Thus, the Jacobian consists of entries of the form of $f(x, y) \doteq \frac{x}{(x+y)^2}$. For $x, y \in [\underline{p}, 1]$, we have $|f(x, y)| \leq \frac{1}{4\underline{p}^2}$. This means that the absolute value of the off-diagonal entries of $\nabla \boldsymbol{z}^{\text{scarf}}(\boldsymbol{p})$ are bounded by $\frac{1}{4\underline{p}^2}$, while the diagonal entries are bounded by $\frac{1}{2\underline{p}^2}$. Hence, for all $\boldsymbol{p} \in [\underline{p}, 1]^3$, it holds that $\|\nabla \boldsymbol{z}(\boldsymbol{p})\|_1 \leq \frac{3}{2\underline{p}^2} + \frac{6}{4\underline{p}^2} = \frac{3}{\underline{p}^2}$. Therefore, by the mean value theorem, $\boldsymbol{z}^{\text{scarf}}$ is $\frac{3}{\underline{p}^2}$-Lipschitz-continuous on $[\underline{p}, 1]^3$, i.e., for all

$\boldsymbol{p}, \boldsymbol{q} \in [\underline{p}, 1]^3$, $\|\boldsymbol{z}(\boldsymbol{p}) - \boldsymbol{z}(\boldsymbol{q})\| \leq \frac{3}{\underline{p}^2} \|\boldsymbol{p} - \boldsymbol{q}\|$. Now, since h is 1-strongly-convex, we have, for all $\boldsymbol{p}, \boldsymbol{q} \in \mathbb{R}^3_+$, $\frac{1}{2} \|\boldsymbol{p} - \boldsymbol{q}\|^2 \leq \mathrm{div}_h(\boldsymbol{p}, \boldsymbol{q})$. Hence, for all $\boldsymbol{p}, \boldsymbol{q} \in [\underline{p}, 1]^3$,

$$\begin{aligned} {}^1\!/\!{}_2 \|\boldsymbol{z}(\boldsymbol{p}) - \boldsymbol{z}(\boldsymbol{q})\|^2 &\leq \frac{1}{2} \left(\frac{3}{\underline{p}^2} \right)^2 \|\boldsymbol{p} - \boldsymbol{q}\|^2 \\ &\leq \left(\frac{3}{\underline{p}^2} \right)^2 \mathrm{div}_h(\boldsymbol{p}, \boldsymbol{q}) \end{aligned}$$

With the above lemma in hand, we can prove the convergence of mirror *extratâtonnement* in the Scarf economy.

Corollary 5.4.4 [Convergence of Mirror *Extrâtonnement* in the Scarf Economy]**.**
Let $\underline{p} \in (0, 1)$. Consider the mirror *extrâtonnement* process run on the Scarf economy $\boldsymbol{z}^{\mathrm{scarf}}$, with a 1-strongly-convex and κ-Lipschitz-smooth kernel function h, any time horizon $t \dot\in \mathbb{N}$, any step size $\eta \dot\in (0, \frac{1}{\sqrt{2}\lambda}]$, a price space $\mathcal{P} \doteq [\underline{p}, 1]^3$, and any initial price vector $\boldsymbol{p}^{(0)} \in \mathcal{P}$. If $\{\boldsymbol{p}^{(t)}, \boldsymbol{p}^{(t+0.5)}\}_t$ is the output price sequence, then $\lim_{t \to \infty} \boldsymbol{p}^{(t+0.5)} = \lim_{t \to \infty} \boldsymbol{p}^{(t)} = \boldsymbol{p}^*$ is a Walrasian equilibrium.

Indeed, mirror *extrâtonnement* is a discrete-time natural price adjustment process that converges to the unique Walrasian equilibrium in the Scarf economy.

5.4.3 Experiments for Mirror *Extratâtonnement* Process

In this section, we first apply the *tâtonnement* and mirror *extratâtonnement* processes with the kernel function $h(\boldsymbol{p}) \doteq \|\boldsymbol{p}\|^2$ to solve the Scarf economy, with the goal of illustrating the differing convergence behaviors between the two price-adjustment processes. We then apply the mirror *extratâtonnement* process with kernel function $h(\boldsymbol{p}) \doteq \|\boldsymbol{p}\|^2$ to solve a number of Arrow-Debreu exchange economies (Arrow and Debreu, 1954), with the goal of demonstrating that our pathwise Bregman-continuity assumption holds, and that mirror *extratâtonnement* can efficiently solve very large Walrasian economies in practice.

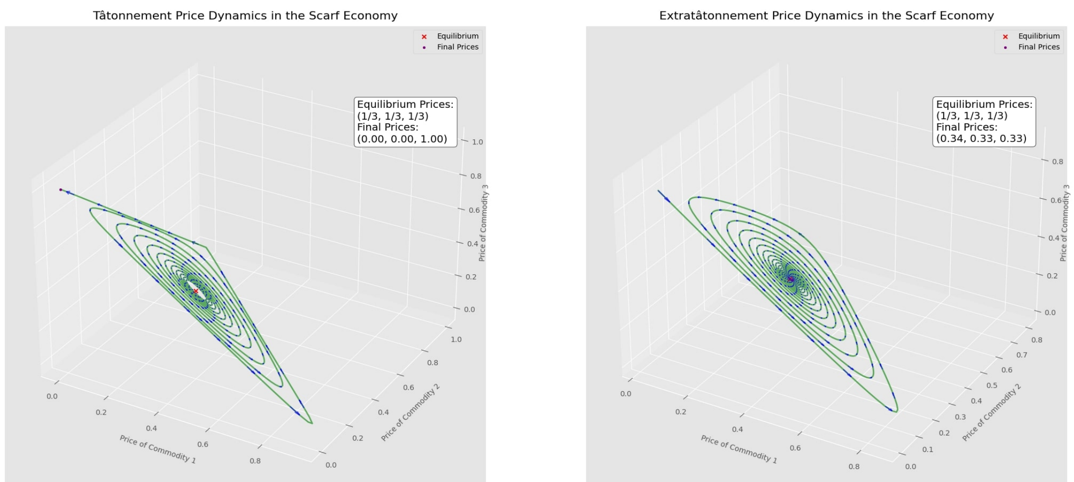

Figure 5.1: Phase Portraits of *Tâtonnement* and *Extratâtonnement* for the Scarf Economy

We record in Figure 5.1 the movement of prices in the Scarf economy for the *tâtonnement* and mirror *extratâtonnement* processes, respectively. As is well-established by now, the price sequence generated by *tâtonnement*, despite starting very close to the equilibrium prices $(1/3, 1/3, 1/3)$ spirals away from those prices, converging to $(0, 0, 1)$, which is *not* a Walrasian equilibrium. In contrast, the prices generated by the mirror *extratâtonnement* process spiral inwards towards the equilibrium prices, despite starting far away from them.

An intuitive explanation of the observed behavior is as follows: As noted above, the continuous-time variant of *tâtonnement* is known to cycle around the equilibrium prices (Scarf, 1960). One way to interpret the discrete-time *tâtonnement* (respectively, mirror *extratâtonnement*) process is as an explicit (respectively, implicit) discretization (Butcher, 2008) of the continuous-time *tâtonnement* dynamics. A well-known fact is that explicit (respectively, implicit) discretization methods are unstable (respectively, stable) when continuous-time dynamics cycle, thus explaining the observed behavior.

An Arrow-Debreu exchange economy $(n, m, \mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ comprises $m \in \mathbb{N}$ commodities, $n \in \mathbb{N}$ consumers each $i \in [n]$ with a consumption space $\mathcal{X}_i$, an endowment of commodities $\boldsymbol{e}_i \in \mathbb{R}^m_+$, and a utility function $u_i : \mathcal{X}_i \to \mathbb{R}$. An Arrow-Debreu exchange economy $(n, m, \mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ can be represented as a bounded continuous competitive economy $(m, \mathcal{Z})$, where the excess demand correspondence is given as $\mathcal{Z}(\boldsymbol{p}) \doteq \sum_{i \in [n]} \underset{\boldsymbol{x}_i \in \mathcal{X}_i : \boldsymbol{x}_i \cdot \boldsymbol{p} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}}{\arg\max} u_i(\boldsymbol{x}_i) - \sum_{i \in [n]} \boldsymbol{e}_i$.[7]

Table 5.1: Summary of Setups for Arrow-Debreu Exchange Economy Experiments

Exp No.	Num. Comm.	Num. Linear Cons.	Num. Cobb -Doug. Cons.	Num. CES $\rho \in (0, 1)$ Cons.	Num. CES $\rho < 0$ Cons.	Num. Leont. Cons.
1	500	0	0	0	0	600
2	500	0	0	0	600	0
3	500	0	0	600	0	0
4	500	0	600	0	0	0
5	500	600	0	0	0	0
6	1000	200	200	200	200	200
7	1000	0	200	200	200	200

We consider the following utility function classes in our experiments: 1. linear: $u_i(\boldsymbol{x}_i) = \sum_{j \in [m]} v_{ij} x_{ij}$; 2. Cobb-Douglas: $u_i(\boldsymbol{x}_i) = \prod_{j \in [m]} x_{ij}^{v_{ij}}$; 3. Leontief: $u_i(\boldsymbol{x}_i) = \min_{j \in [m]} \{ x_{ij}/v_{ij} \}$; and 4. CES: $u_i(\boldsymbol{x}_i) = \sqrt[\rho_i]{\sum_{j \in [m]} v_{ij} x_{ij}^{\rho_i}}$. Each class of utility functions is parameterized by

[7]We refer the reader to Chapter 10 on additional background and definitions regarding Arrow-Debreu exchange economies.

a vector of valuations $\boldsymbol{v}_i \in \mathbb{R}^n_+$, where each v_{ij} quantifies the value of commodity j to consumer i. We summarize the experiments we run in Table 5.1. The parameters of each economy are initialized randomly according to the uniform random distribution.[8] We record the results of our experiments in Figure 5.2, describing for what values of $\varepsilon \geq 0$ the prices generated throughout the algorithm converge to an ε-Walrasian equilibrium.

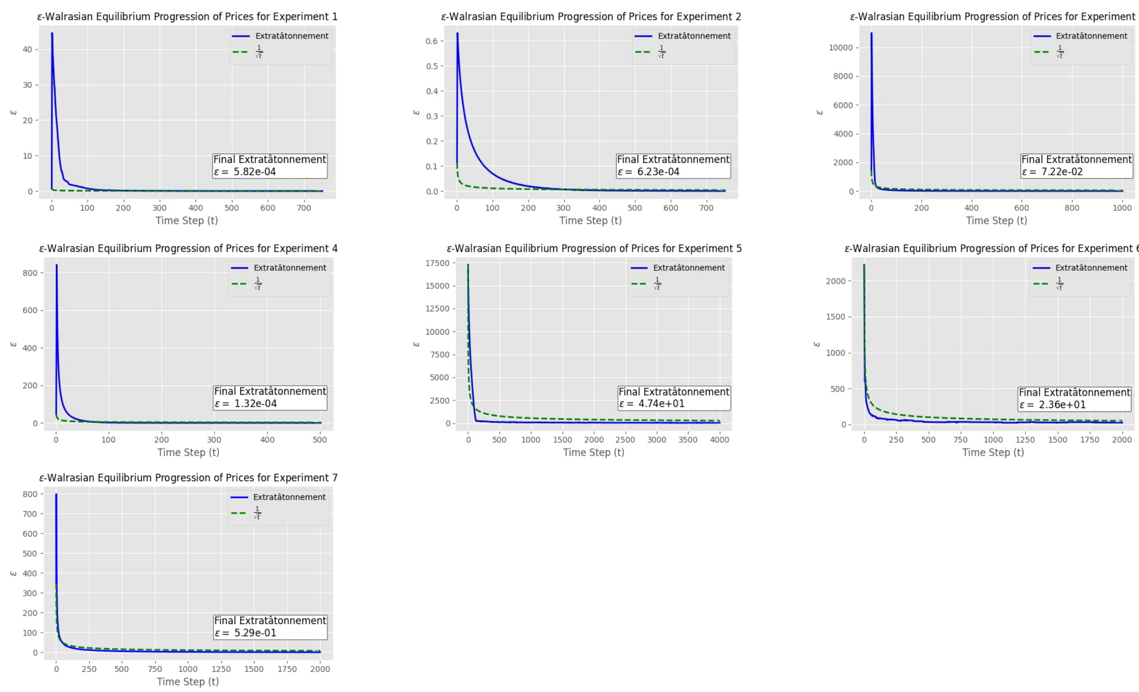

Figure 5.2: Results of Experiments 1-7

We observe that in all our experiments except Experiments 5 and 6—which include linear consumers and are not as such covered by our theory, as the excess demand in such economies is not singleton-valued—the mirror *extratâtonnement* process converges to a Walrasian equilibrium. In all experiments, we verify and confirm that pathwise Bregman-continuity holds, thus justifying this assumption. Finally, while our experiments obey our theory, which suggests best-iterate convergence to an ε-Walrasian equilibrium in $1/\varepsilon^2$ time-steps, we observe last-iterate convergence in Experiment 4, corresponding to the case of Cobb-Douglas consumers, for which even *tâtonnement* is known to converge in last

[8] For reproducibility purposes, we include our code ready to run on https://github.com/denizalp/extratatonnement, and include all details of our experimental setup in Section 5.4.3.

iterates. This finding suggests that convergence in last iterates might not be possible with the mirror *extratâtonnement* process.

5.5 Merit Function Methods for Walrasian Equilibrium

5.5.1 Merit Function Minimization via Second-Order Price-Adjustment Process

Going beyond Minty Arrow-Debreu economies, we can also apply the merit function method we introduced in Chapter 4 to solve general Walrasian economies with excess demand functions that are sufficiently smooth. Recall that by Theorem 5.2.1 the set of Walrasian equilibria of any Walrasian economy $(m, \boldsymbol{z})$ is equal to the set of strong solutions $\mathcal{SVI}(\mathbb{R}^m_+, -\boldsymbol{z})$ of the VI $(\mathbb{R}^m_+, -\boldsymbol{z})$. As such, by applying the merit function derived for VIs in Lemma 4.4.1, we define the following merit function for Walrasian equilibrium:

$$\Xi_\alpha(\boldsymbol{p}) \doteq \max_{\boldsymbol{q}\in\mathbb{R}^m_+} \langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{q} - \boldsymbol{p}\rangle - \frac{\alpha}{2}\left\|\boldsymbol{q} - \boldsymbol{p}\right\|^2 \tag{5.12}$$

We have the following corollary of Lemma 4.4.1, which characterizes this merit function.

Corollary 5.5.1 [Merit function for Walrasian equilibrium]**.**
Given a Walrasian economy $(m, \boldsymbol{z})$, for any $\alpha \geq 0$, the set of Walrasian equilibria of $(m, \mathcal{Z})$ is equal to $\arg\min_{\boldsymbol{p}\in\mathbb{R}^m_+} \Xi_\alpha(\boldsymbol{p})$. Furthermore, if $\alpha > 0$, then $\arg\max_{\boldsymbol{q}\in\mathbb{R}^m_+}\{\langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{q} - \boldsymbol{p}\rangle - \frac{\alpha}{2}\left\|\boldsymbol{q} - \boldsymbol{p}\right\|^2\} = \{\boldsymbol{q}^*(\boldsymbol{p})\}$, where

$$\boldsymbol{q}^*(\boldsymbol{p}) = \arg\max_{\boldsymbol{q}\in\mathbb{R}^m_+} \langle \boldsymbol{z}(\boldsymbol{p}), \boldsymbol{p} - \boldsymbol{q}\rangle - \frac{\alpha}{2}\left\|\boldsymbol{q} - \boldsymbol{p}\right\|^2 = \Pi_{\mathbb{R}^m_+}\left[\boldsymbol{p} + \frac{1}{\alpha}\boldsymbol{z}(\boldsymbol{p})\right] \quad .$$

In addition, Ξ_α can be expressed as follows:

$$\Xi_\alpha(\boldsymbol{p}) = \max_{\boldsymbol{q}\in\mathbb{R}^m_+} \frac{\alpha}{2}\left[\left\|\frac{1}{\alpha}\boldsymbol{z}(\boldsymbol{p})\right\|^2 - \left\|\boldsymbol{q} - \left(\boldsymbol{p} - \frac{1}{\alpha}\boldsymbol{f}(\boldsymbol{x})\right)\right\|^2\right] \quad ,$$

with its gradient given by:

$$\nabla\Xi_\alpha(\boldsymbol{p}) \doteq \boldsymbol{z}(\boldsymbol{p}) - \left(\nabla\boldsymbol{z}(\boldsymbol{p}) + \alpha\mathbb{I}\right)\left(\boldsymbol{q}^*(\boldsymbol{p}) - \boldsymbol{p}\right) \quad .$$

5.5.2 Mirror Potential Algorithm for Walrasian Economies

With the above lemma in hand, we can minimize Ξ_α using the mirror potential algorithm (Algorithm 5). Note that like the mirror potential algorithm, the algorithm that arises is a second-order price adjustment process. The following corollary is obtained by applying Theorem 4.4.1 to the regularized primal gap as defined in Equation (5.12).

Theorem 5.5.1 [Mirror potential algorithm for Walrasian equilibrium]**.**

Consider a Walrasian economy $(m, \boldsymbol{z})$ with an excess demand function $\boldsymbol{z}$ that is λ-Lipschitz-continuous and β-Lipschitz-smooth, a 1-strongly-convex kernel function h, $\alpha \geq 0$, $\eta \in \left(0, \frac{1}{2(2\beta\alpha \mathrm{diam}(\mathcal{X})^2+1+2\lambda)}\right]$, and $\boldsymbol{x}^{(0)} \in \mathcal{X}$.

Consider the mirror potential algorithm (Algorithm 5) run with the regularized primal gap Ξ_α as defined in Equation (5.12), the kernel function h, an arbitrary time horizon $\tau \in \mathbb{N}$, the step size η, and the initial iterate $\boldsymbol{x}^{(0)}$. The output sequence $\{\boldsymbol{x}^{(t)}\}_t$ satisfies the following convergence bound to a stationary point of Ξ_α:

$$\min_{k=0,1,\ldots,\tau-1} \max_{\boldsymbol{x}\in\mathcal{X}} \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x}\rangle \leq \frac{2\Xi_\alpha(\boldsymbol{x}^{(0)})}{\tau}$$

In addition, if $\boldsymbol{x}^{(\tau)}_{\text{best}} \in \arg\min_{\boldsymbol{x}^{(k)}:k=0,\ldots,\tau-1} \max_{\boldsymbol{x}\in\mathcal{X}} \langle \nabla \Xi_\alpha(\boldsymbol{x}^{(k)}), \boldsymbol{x}^{(k)} - \boldsymbol{x}\rangle$, then for some $\tau \in O(1/\varepsilon)$, $\boldsymbol{x}^{(\tau)}_{\text{best}}$ is an ε-stationary point of Ξ_α.

Remark 5.5.1 [Walrasian equilibrium under the law of demand and supply]**.**

As pointed out in Remark 4.4.1, when the excess demand is monotone, stationary points of the regularized primal gap function are global solutions of the economy. In the case of Walrasian economies, monotonicity of the excess demand is known as the law of supply and demand (see, Definition 5.4.8). As such, to the best of our knowledge, Theorem 5.5.1 is the first polynomial-time guarantee for Walrasian equilibrium in general Walrasian economies (i.e., beyond Arrow-Debreu economies), whose excess demand satisfies the law of demand and supply.

Remark 5.5.2 [On Lipschitz-continuity and smoothness]**.**

Recall that we can ensure that the excess demand is Lipschitz-continuous by Lemma 5.4.2, assuming that the price elasticity of the excess demand is bounded. Nevertheless, to the best of our knowledge, there are no known settings of the market parameters that obtain Lipschitz-smoothness of the excess demand. That said, as the Lipschitz-continuity and Lipschitz-smoothness of a function can be approximated from data, this assumption remains realistic (see, for instance, (Wood and Zhang, 1996)).

Having established very broad results for the convergence of price-adjustment processes for Walrasian economies, in the next chapter, we provide an example of a Walrasian economy, namely Fisher markets (Brainard et al., 2000), which have found a great deal of applications in real-world resource allocation problems. This application will also demonstrate that in more restricted Walrasian economies, using a more fine-grained analysis, the convergence of *tâtonnement* processes can be guaranteed, complementing the results established in this chapter.

Chapter 6

Homothetic Fisher Markets

In this chapter, we turn our attention to identifying a large class of Walrasian economies for which the mirror *tâtonnement process* converges to a Walrasian equilibrium. To this end, we make strides towards analyzing the convergence of discrete-time *tâtonnement* in **homothetic Fisher markets**, i.e., Walrasian economies with a fixed supply, and a demand generated by utility-maximizing consumers whose utility functions are given by continuous and homogeneous functions.[1] An important concept in consumer theory is a buyer's Hicksian demand, i.e., consumptions that minimize expenditure while achieving a desired utility level. We identify the maximum elasticity of the Hicksian demand, i.e., the maximum percentage change in the Hicksian demand of any good w.r.t. the change in the price of some other good, as an economic parameter sufficient to capture and explain a range of convergent and non-convergent *tâtonnement* behaviors in a broad class of markets. In particular, we prove the convergence of *tâtonnement* in homothetic Fisher markets with bounded elasticity of Hicksian demand, i.e., Fisher markets in which consumers have preferences represented by homogeneous utility functions for which the elasticity of their Hicksian demand is bounded.

[1]We refer to Fisher markets that comprise buyers with a certain utility function by the name of the utility function, e.g., we call a Fisher market that comprises buyers with Leontief utility functions a Leontief Fisher market. We omit the "continuous" qualifier as Walrasian equilibrium is not guaranteed to exist when utilities are not continuous.

6.1 Background

6.1.1 Mirror Descent

Consider the optimization problem $\min_{\boldsymbol{x}\in V} f(\boldsymbol{x})$, where $f : \mathbb{R}^n \to \mathbb{R}$ is a differentiable convex function and V is the feasible set of solutions. A standard method for solving this problem is the **mirror descent algorithm** (Boyd et al., 2004):

$$\boldsymbol{x}(t+1) = \arg\min_{\boldsymbol{x}\in V} \{\ell_f(\boldsymbol{x}, \boldsymbol{x}(t)) + \gamma_t \mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}(t))\} \qquad \text{for } t = 0, 1, 2, \ldots \tag{6.1}$$

$$\boldsymbol{x}(0) \in \mathbb{R}^n \tag{6.2}$$

Here, $\gamma_t > 0$ is the step size at time t, $\ell_f(\boldsymbol{x}, \boldsymbol{y})$ is the **linear approximation** of f at $\boldsymbol{y}$, that is $\ell_f(\boldsymbol{x}, \boldsymbol{y}) = f(\boldsymbol{y}) + \nabla f(\boldsymbol{y})^T(\boldsymbol{x} - \boldsymbol{y})$, and $\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{x}(t))$ is the **Bregman divergence** of a convex differentiable **kernel** function $h(\boldsymbol{x})$ defined as $\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{y}) = h(\boldsymbol{x}) - \ell_h(\boldsymbol{x}, \boldsymbol{y})$ (Bregman, 1967). In particular, when $h(\boldsymbol{x}) = \frac{1}{2}||\boldsymbol{x}||_2^2$, $\mathrm{div}_h(\boldsymbol{x}, \boldsymbol{y}) = \frac{1}{2}||\boldsymbol{x} - \boldsymbol{y}||_2^2$. In this case, mirror descent reduces to projected gradient descent (Boyd et al., 2004). If instead the kernel is the weighted entropy $h(\boldsymbol{x}) = \sum_{i\in[n]} (x_i \log(x_i) - x_i)$, the Bregman divergence reduces to the **generalized Kullback-Leibler (KL) divergence** (Joyce, 2011):

$$\mathrm{div}_{\mathrm{KL}}(\boldsymbol{x}, \boldsymbol{y}) = \sum_{i\in[n]} \left[x_i \log\left(\frac{x_i}{y_i}\right) - x_i + y_i \right] \ , \tag{6.3}$$

which, when $V = \mathbb{R}^m_{++}$, yields the following simplified **entropic descent** update rule:

$$\forall j \in [m] \qquad x_j^{(t+1)} = x_j^{(t)} \exp\left\{ \frac{-\mathcal{D}_{x_j} f(\boldsymbol{x}^{(t)})}{\gamma_t} \right\} \qquad \text{for } t = 0, 1, 2, \ldots \tag{6.4}$$

$$x_j^{(0)} \in \mathbb{R}_{++} \tag{6.5}$$

A function f is said to be **γ-Bregman-smooth** (Cheung et al., 2018) w.r.t. a Bregman divergence with kernel function h if $f(\boldsymbol{x}) \le \ell_f(\boldsymbol{x}, \boldsymbol{y}) + \gamma \mathrm{div}_h(\boldsymbol{x}, \boldsymbol{y})$. Birnbaum et al. (2011) showed that if the objective function $f(\boldsymbol{x})$ of a convex optimization problem is **γ-Bregman** w.r.t. to some Bregman divergence div_h, then mirror descent with Bregman divergence div_h converges to an optimal solution $f(\boldsymbol{x}^*)$ at a rate of $O(1/t)$. We require a slightly modified version of this theorem, introduced by Cheung et al. (2013), where it suffices for the γ-Bregman-smoothness property to hold only for consecutive pairs of iterates.

Theorem 6.1.1 [Birnbaum et al. (2011),Cheung et al. (2013)].
Let $\{\boldsymbol{x}^t\}_t$ be the iterates generated by mirror descent with Bregman divergence div_h. Suppose f and h are convex, and for all $t \in \mathbb{N}$ and for some $\gamma > 0$, it holds that $f(\boldsymbol{x}^{(t+1)}) \leq \ell_f(\boldsymbol{x}^{(t+1)}, \boldsymbol{x}^{(t)}) + \gamma \mathrm{div}_h(\boldsymbol{x}^{(t+1)}, \boldsymbol{x}^{(t)})$. If $\boldsymbol{x}^*$ is a minimizer of f, then the following holds for mirror descent with fixed step size γ: for all $t \in \mathbb{N}$, $f(\boldsymbol{x}^{(t)}) - f(\boldsymbol{x}^*) \leq \gamma/t \, \mathrm{div}_h(\boldsymbol{x}^*, \boldsymbol{x}^{(0)})$.

6.1.2 Consumer Theory Primer

Let $\mathcal{X} = \mathbb{R}^m_+$ be a set of possible consumptions over m goods s.t. for any $\boldsymbol{x} \in \mathbb{R}^m_+$ and $j \in [m]$, $x_j \geq 0$ represents the amount of good $j \in [m]$ consumed by consumer (hereafter, buyer) i. The preferences of buyer i over different consumptions of goods can be represented by a **preference relation** $\succeq_i$ over $\mathcal{X}$ such that the buyer (resp. weakly) prefers a choice $\boldsymbol{x} \in \mathcal{X}$ to another choice $\boldsymbol{y} \in \mathcal{X}$ iff $\boldsymbol{x} \succ_i \boldsymbol{y}$ (resp. $\boldsymbol{x} \succeq_i \boldsymbol{y}$). A preference relation is said to be **complete** iff for all $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{X}$, either $\boldsymbol{x} \succeq_i \boldsymbol{y}$ or $\boldsymbol{y} \succeq_i \boldsymbol{x}$, or both. A preference relation is said to be **transitive** if, for all $\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z} \in \mathcal{X}$, $\boldsymbol{x} \succeq_i \boldsymbol{z}$ whenever $\boldsymbol{x} \succeq_i \boldsymbol{y}$ and $\boldsymbol{y} \succeq_i \boldsymbol{z}$. A preference relation is said to be **continuous** if for any sequence $\{\boldsymbol{x}^{(n)}, \boldsymbol{y}^{(n)}\}_{n \in \mathbb{N}_+} \subset \mathcal{X} \times \mathcal{X}$ $(\boldsymbol{x}^{(n)}, \boldsymbol{y}^{(n)}) \to (\boldsymbol{x}, \boldsymbol{y})$ and $\boldsymbol{x}^{(n)} \succeq_i \boldsymbol{y}^{(n)}$ for all $n \in \mathbb{N}_+$, it also holds that $\boldsymbol{x} \succeq_i \boldsymbol{y}$. A preference relation $\succeq_i$ is said to be **locally non-satiated** iff for all $\boldsymbol{x} \in \mathcal{X}$ and $\epsilon > 0$, there exists $\boldsymbol{y} \in \mathcal{B}_\epsilon(\boldsymbol{x})$ such that $\boldsymbol{y} \succ_i \boldsymbol{x}$. A utility function $u_i : \mathcal{X} \to \mathbb{R}_+$ assigns a positive real value[2] to elements of $\mathcal{X}$, i.e., to all possible consumptions. Every continuous utility function represents some complete, transitive, and continuous preference relation $\succeq_i$ over goods s.t. if $u_i(\boldsymbol{x}) \geq u_i(\boldsymbol{y})$ for two bundles of goods $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^m$, then $\boldsymbol{x} \succeq_i \boldsymbol{y}$ (Arrow et al., 1971).

In this paper, we consider the general class of **homothetic** preferences $\succeq_i$ s.t. for any consumption $\boldsymbol{x}, \boldsymbol{y} \in \boldsymbol{X}$ and $\lambda \in \mathbb{R}_+$, $\boldsymbol{x} \succeq_i \boldsymbol{y}$ and $\boldsymbol{y} \succeq_i \boldsymbol{x}$ implies $\lambda\boldsymbol{x} \succeq_i \lambda\boldsymbol{y}$ and $\lambda\boldsymbol{y} \succeq_i \lambda\boldsymbol{x}$, respectively. A preference relation $\succeq_i$ is complete, transitive, continuous, and homothetic

[2] Without loss of generality, we assume that utility functions are positive real-valued functions, since any real-valued function can be made positive real-valued by passing it through the monotonic transformation $\boldsymbol{x} \mapsto e^{\boldsymbol{x}}$ without affecting the underlying preference relation.

iff it can be represented via a continuous and homogeneous utility function u_i of arbitrary degree (Arrow et al., 1971).[3] We note that any homogeneous utility function u_i represents locally non-satiated preferences, since for all $\epsilon > 0$ and $\boldsymbol{x} \in \mathcal{X}$, there exists an allocation $(1 + \varepsilon/\|\boldsymbol{x}\|)\boldsymbol{x}$ s.t. $u_i((1 + \varepsilon/\|\boldsymbol{x}\|)\boldsymbol{x}) = (1 + \varepsilon/\|\boldsymbol{x}\|)u_i(\boldsymbol{x}) > u_i(\boldsymbol{x})$, and $[\boldsymbol{x} - (1 + \varepsilon/\|\boldsymbol{x}\|)\boldsymbol{x}] \in \mathcal{B}_\varepsilon(\boldsymbol{x})$.

The class of homogeneous utility functions includes the well-known **constant elasticity of substitution (CES)** utility function family, parameterized by a substitution parameter $-\infty \leq \rho_i \leq 1$, and given by $u_i(\boldsymbol{x}_i) = \sqrt[\rho_i]{\sum_{j \in [m]} v_{ij} x_{ij}^{\rho_i}}$ with each utility function parameterized by the vector of valuations $\boldsymbol{v}_i \in \mathbb{R}_+^n$, where each v_{ij} quantifies the value of good j to buyer i. CES utilities are said to be **gross substitutes** (resp. **gross complements**) CES if $\rho_i > 0$ ($\rho_i < 0$). **Linear utility** functions are obtained when ρ is 1 (goods are perfect substitutes), while **Cobb-Douglas** and **Leontief utility** functions are obtained when $\rho \to 0$ and $\rho \to -\infty$ (goods are perfect complements), respectively:

Linear:	Cobb-Doulas:	Leontief:
$u_i(\boldsymbol{x}_i) = \sum_{j \in [m]} v_{ij} x_{ij}$	$u_i(\boldsymbol{x}_i) = \prod_{j \in [m]} x_{ij}^{v_{ij}}$	$u_i(\boldsymbol{x}_i) = \min_{j: v_{ij} \neq 0} \frac{x_{ij}}{v_{ij}}$

Associated with any consumption $\boldsymbol{x} \in \mathcal{X}$ are **prices** $\boldsymbol{p} \in \mathbb{R}_+^m$ s.t. for all goods $j \in [m]$, $p_j \geq 0$ denotes the price of good j. A **demand correspondence** $\mathcal{F} : \mathbb{R}_+^m \to \mathcal{X}$ takes as input prices $\boldsymbol{p} \in \mathbb{R}_+^m$ and outputs a set of consumptions $\mathcal{F}(\boldsymbol{p})$. If $\mathcal{F}$ is singleton-valued for all $\boldsymbol{p} \in \mathbb{R}_+^m$, then it is called a **demand function**. Given a demand function $\boldsymbol{f}$, we define the **elasticity** $\epsilon_{f_i, x_j} : \mathbb{R}^m \to \mathbb{R}$ of output $f_i(\boldsymbol{x})$ w.r.t. the jth input x_j evaluated at $\boldsymbol{x} = \boldsymbol{y}$ as $\epsilon_{f_i, x_j}(\boldsymbol{y}) = \mathcal{D}_{x_j} f_i(\boldsymbol{y}) \frac{y_j}{f_i(\boldsymbol{y})}$.

A good $j \in [m]$ is said to be a **substitute (resp. complement) w.r.t. a demand function** $\boldsymbol{f}$ for a good $k \in [m] \setminus \{j\}$ if the demand $f_j(\boldsymbol{p})$ is increasing (resp. decreasing) in p_k. If a buyer's demand $f_j(\boldsymbol{p})$ for good j is instead weakly increasing (resp. decreasing), good j is said to be a **weak substitute** (resp. **weak complement**) for good k.

[3] Throughout this work, without loss of generality, we assume that complete, transitive, continuous, and homothetic preference relations are represented via a homogeneous utility function of degree 1, since any homogeneous utility function of degree k can be made homogeneous of degree 1 without affecting the underlying preference relation by passing the utility function through the monotonic transformation $x \mapsto \sqrt[k]{x}$.

Next, we define the **consumer functions** (Mas-Colell et al., 1995; Jehle, 2001). The **indirect utility function** $v_i : \mathbb{R}_+^m \times \mathbb{R}_+ \to \mathbb{R}_+$ takes as input prices $\boldsymbol{p}$ and a budget b_i and outputs the maximum utility the buyer can achieve at that prices within that budget, i.e., $v_i(\boldsymbol{p}, b_i) = \max_{\boldsymbol{x} \in \mathcal{X} : \boldsymbol{p} \cdot \boldsymbol{x} \leq b_i} u_i(\boldsymbol{x})$.

The **Marshallian demand** is a correspondence $\mathcal{D}_i : \mathbb{R}_+^m \times \mathbb{R}_+ \rightrightarrows \mathcal{X}$ that takes as input prices $\boldsymbol{p}$ and a budget b_i and outputs the utility-maximizing allocations of goods at that budget, i.e., $\mathcal{D}_i(\boldsymbol{p}, b_i) = \arg\max_{\boldsymbol{x} \in \mathcal{X} : \boldsymbol{p} \cdot \boldsymbol{x} \leq b_i} u_i(\boldsymbol{x})$.

The **expenditure function** $e_i : \mathbb{R}_+^m \times \mathbb{R}_+ \to \mathbb{R}_+$ takes as input prices $\boldsymbol{p}$ and a utility level ν_i and outputs the minimum amount the buyer must spend to achieve that utility level at those prices, i.e., $e_i(\boldsymbol{p}, \nu_i) = \min_{\boldsymbol{x} \in \mathcal{X} : u_i(\boldsymbol{x}) \geq \nu_i} \boldsymbol{p} \cdot \boldsymbol{x}$. If the utility function u_i is continuous, then the expenditure function is continuous and homogeneous of degree 1 in $\boldsymbol{p}$ and ν_i jointly, non-decreasing in $\boldsymbol{p}$, strictly increasing in ν_i, and concave in $\boldsymbol{p}$.

The **Hicksian demand** is a correspondence $\mathcal{H}_i : \mathbb{R}_+^m \times \mathbb{R}_+ \rightrightarrows \mathbb{R}_+$ that takes as input prices $\boldsymbol{p}$ and a utility level ν_i and outputs the cost-minimizing allocations of goods at those prices and utility level, i.e., $\mathcal{H}_i(\boldsymbol{p}, \nu_i) = \arg\min_{\boldsymbol{x} \in \mathcal{X} : u_i(\boldsymbol{x}) \geq \nu_i} \boldsymbol{p} \cdot \boldsymbol{x}$.

6.1.3 Fisher Markets

A **Fisher market** $(n, m, \boldsymbol{u}, \boldsymbol{b})$, denoted $(\boldsymbol{u}, \boldsymbol{b})$ when clear from context, consists of n traders and m goods. A **Fisher market** consists of n buyers and m divisible goods (Brainard et al., 2000). Each buyer $i \in [n]$ has a budget $b_i \in \mathbb{R}_{++}$ and a utility function $u_i : \mathbb{R}_+^m \to \mathbb{R}$. As is standard in the literature, we assume there is one unit of each good, and one unit of currency available in the market, i.e. $\sum_{i \in [n]} b_i = 1$ (Nisan and Roughgarden, 2007). We denote $\boldsymbol{u} = (u_1, \ldots, u_n)$, and $\boldsymbol{b} \doteq (b_1, \ldots, b_n)$.

An **allocation** $\boldsymbol{X}$ is a map from goods to buyers, represented as a matrix s.t. $x_{ij} \geq 0$ denotes the amount of good $j \in [m]$ allocated to buyer $i \in [n]$. Goods are assigned **prices** $\boldsymbol{p} \in \mathbb{R}_+^m$.

When the buyers' utility functions in a Fisher market are all of the same type, we qualify the market by the name of the utility function, e.g., a linear Fisher market. A **mixed CES Fisher market** is a Fisher market which comprises CES buyers with possibly different substitution parameters. Considering properties of goods, rather than buyers, a (Fisher) market satisfies **gross substitutes** (resp. **gross complements**) if all pairs of goods in the market are gross substitutes (resp. gross complements). We define net substitute Fisher markets and net complements Fisher markets similarly. We refer the reader to Figure 3.1a for a summary of the relationships among various Fisher markets.

Definition 6.1.1 [Fisher Equilibrium]**.**
A tuple $(\boldsymbol{X}^*, \boldsymbol{p}^*)$ is said to be a **Fisher equilibrium** of a Fisher market $(\boldsymbol{u}, \boldsymbol{b})$ iff

(Utility maximization)	Buyers maximize their utility constrained by their budget, i.e., $\forall i \in [n], \boldsymbol{x}_i^* \in \mathcal{D}_i(\boldsymbol{p}^*, b_i)$;
(Feasibility)	$\forall j \in [m], \sum_{i \in [n]} x_{ij}^* \leq 1$
(Walras' law)	$\boldsymbol{p}^* \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i^* - \mathbf{1}_m\right)$

Any Fisher market $(\boldsymbol{u}, \boldsymbol{b})$ can be represented as a Walrasian economy. To this end, overloading notation, define the **aggregate demand correspondence** $\mathcal{D} : \mathbb{R}_+^m \rightrightarrows \mathbb{R}_+^m$ at prices $\boldsymbol{p}$ as the sum of the Marshallian demand at $\boldsymbol{p}$, given budgets $\boldsymbol{b}$, i.e., $\mathcal{D}(\boldsymbol{p}) = \sum_{i \in [n]} \mathcal{D}_i(\boldsymbol{p}, b_i)$. The **excess demand correspondence** $\mathcal{Z} : \mathbb{R}^m \rightrightarrows \mathbb{R}^m$ of a Fisher market $(\boldsymbol{u}, \boldsymbol{b})$, which takes as input prices and outputs a set of excess demands at those prices, Any Fisher market $(\boldsymbol{u}, \boldsymbol{b})$ can be represented as a Walrasian economy $(m, \mathbb{R}_+^m, \mathcal{Z})$ where $\mathcal{Z}$ is defined as the difference between the aggregate demand for and the supply of each good: i.e., $\mathcal{Z}(\boldsymbol{p}) = \mathcal{D}(\boldsymbol{p}) - \mathbf{1}_m$ where $\mathbf{1}_m$ is the vector of ones of size m, and $\mathcal{D}(\boldsymbol{p}) - \mathbf{1}_m = \{\boldsymbol{x} - \mathbf{1}_m \mid \forall \boldsymbol{x} \in \mathcal{D}(\boldsymbol{p})\}$. Note that the the set of Fisher equilibrium prices of any Fisher market $(\boldsymbol{u}, \boldsymbol{b})$ is equal to the set of Walrasian equilibria of $(m, \mathcal{Z})$.

6.2 Homothetic Fisher Markets

We now turn our attention to the computation of Fisher equilibrium in Fisher markets. We will restrict ourselves to a large class of Fisher markets, namely homothetic Fisher markets.

Definition 6.2.1 [Homothetic Fisher Markets]**.**
A **homothetic Fisher market** is a Fisher market $(\boldsymbol{u}, \boldsymbol{b})$ s.t. for each buyer $i \in [n]$:

$$\begin{aligned} &\text{(Continuity)} && u_i \text{ is continuous;} \\ &\text{(Homothetic preferences)} && u_i \text{ is homogeneous, i.e., for all } \lambda \geq 0,\ \lambda u_i(\boldsymbol{x}_i) = u_i(\lambda \boldsymbol{x}_i). \end{aligned}$$

Suppose that $(\boldsymbol{u}, \boldsymbol{b})$ is a continuous, concave, and homogeneous Fisher market. The optimal solutions $(\boldsymbol{X}^*, \boldsymbol{p}^*)$ to the primal and dual of **Eisenberg-Gale program** constitute a Fisher equilibrium of $(\boldsymbol{u}, \boldsymbol{b})$ (Devanur et al., 2002; Eisenberg and Gale, 1959; Jain et al., 2005):[4]

Primal

$$\begin{aligned} \max_{\boldsymbol{X} \in \mathbb{R}_+^{n \times m}} \quad & \sum_{i \in [n]} b_i \log\left(u_i(\boldsymbol{x}_i)\right) \\ \text{subject to} \quad & \sum_{i \in [n]} x_{ij} \leq 1 \qquad \forall j \in [m] \end{aligned}$$

Dual

$$\min_{\boldsymbol{p} \in \Delta_m} \quad \sum_{j \in [m]} p_j + \sum_{i \in [n]} \left[b_i \log\left(v_i(\boldsymbol{p}, b_i)\right) - b_i\right]$$

We now introduce a convex program which is equivalent to the Eisenberg-Gale convex program, but whose optimal value differs from that of the Eisenberg-Gale convex program by an additive constant. Before presenting our program, we present several preliminary lemmas. All omitted proofs can be found in Section 7.2.

The next lemma establishes an important property of the indirect utility and expenditure functions in CCH Fisher markets that we heavily exploit in this work, namely that the derivative of the indirect utility function with respect to b_i—the bang-per-buck—is constant across all budget levels. Likewise, the derivative of the expenditure function with respect to ν_i—the buck-per-bang—is constant across all utility levels. In other words, both functions effectively depend only on prices. Not only are the bang-per-buck and the buck-per-bang

[4]The dual as presented here was formulated by Goktas et al. (2022b).

constant, they equal $v_i(\boldsymbol{p}, 1)$ and $e_i(\boldsymbol{p}, 1)$, respectively, namely their values at exactly one unit of budget and one unit of (indirect) utility.

An important consequence of this lemma is that, by picking prices that maximize a buyer's bang-per-buck, we not only maximize their bang-per-buck at all budget levels, but we further maximize their total indirect utility, given their *known* budget. In particular, given prices $\boldsymbol{p}^*$ that maximize a buyer's bang-per-buck at budget level 1, we can easily calculate the buyer's total (indirect) utility at budget b_i by simply multiplying their bang-per-buck by b_i: i.e., $v_i(\boldsymbol{p}^*, b_i) = b_i v_i(\boldsymbol{p}^*, 1)$. Here, we see quite explicitly the homogeneity assumption at work.

Analogously, by picking prices that maximize a buyer's buck-per-bang, we not only maximize their buck-per-bang at all utility levels, but we further maximize the buyer's total expenditure, given their *unknown* optimal utility level. As above, given prices $\boldsymbol{p}^*$ that minimize a buyer's buck-per-bang at utility level 1, we can easily calculate the buyer's total expenditure at utility level ν_i by simply multiplying their buck-per-bang by ν_i: i.e., $e_i(\boldsymbol{p}^*, \nu_i) = \nu_i e_i(\boldsymbol{p}^*, 1)$.

In sum, solving for optimal prices at any budget level, or analogously at any utility level, requires only a single optimization, in which we solve for optimal prices at budget level, or utility level, 1.

Lemma 6.2.1.

If u_i is continuous and homogeneous of degree 1, then $v_i(\boldsymbol{p}, b_i)$ and $e_i(\boldsymbol{p}, \nu_i)$ are differentiable in b_i and ν_i, resp. Further, $\mathcal{D}_{b_i} v_i(\boldsymbol{p}, b_i) = \{v_i(\boldsymbol{p}, 1)\}$ and $\mathcal{D}_{\nu_i} e_i(\boldsymbol{p}, \nu_i) = \{e_i(\boldsymbol{p}, 1)\}$.

The next lemma provides further insight into why CCH Fisher markets are easier to solve than non-CCH Fisher markets. The lemma states that the bang-per-buck, i.e., the marginal utility of an additional unit of budget, is equal to the inverse of its buck-per-bang, i.e., the marginal cost of an additional unit of utility. Consequently, by setting prices so as to minimize the buck-per-bang of buyers, we can also maximize their bang-per-buck. Since the buck-per-bang is a function of prices only, and not of prices and allocations together, this lemma effectively decouples the calculation of equilibrium prices from the calculation of equilibrium allocations, which greatly simplifies the problem of computing Fisher equilibria in CCH Fisher markets.

Corollary 6.2.1.

If buyer i's utility function u_i is CCH, then

$$\frac{1}{e_i(\boldsymbol{p},1)} = \frac{1}{\frac{\partial e_i(\boldsymbol{p},\nu_i)}{\partial \nu_i}} = \frac{\partial v_i(\boldsymbol{p},b_i)}{\partial b_i} = v_i(\boldsymbol{p},1) \quad . \tag{6.6}$$

We can now present our characterization of the dual of the Eisenberg-Gale program via expenditure functions. While Devanur et al. (Devanur et al., 2016) provided a method to construct a similar program to that given in Theorem 6.2.1 for specific utility functions, their method does not apply to arbitrary CCH utility functions. The proof of this theorem can be found in Section 7.2.

Theorem 6.2.1 [New Convex Program for Homothetic Fisher Markets]**.**

The optimal solutions $(\boldsymbol{X}^*, \boldsymbol{p}^*)$ to the following primal and dual convex programs correspond to Fisher equilibrium allocations and prices, respectively, of the homothetic Fisher market $(\boldsymbol{u}, \boldsymbol{b})$:

Primal

$$\begin{aligned} \max_{\boldsymbol{X} \in \mathbb{R}_+^{n \times m}} \quad & \sum_{i \in [n]} \left[b_i \log u_i \left(\frac{\boldsymbol{x}_i}{b_i} \right) + b_i \right] \\ \text{subject to} \quad & \sum_{i \in [n]} x_{ij} \leq 1 \qquad \forall j \in [m] \end{aligned}$$

Dual

$$\min_{\boldsymbol{p} \in \Delta_m} \psi(\boldsymbol{p}) \doteq \sum_{j \in [m]} p_j - \sum_{i \in [n]} b_i \log \left(e_i(\boldsymbol{p}, 1) \right)$$

Our new convex program for CCH Fisher markets makes plain the duality structure between utility functions and expenditure functions that is used to compute "shadow" prices for allocations. In particular, $e_i(\boldsymbol{p}, \nu_i)$ is the Fenchel conjugate of the indicator function $\chi_{\{\boldsymbol{x}: u_i(\boldsymbol{x}_i) \geq \nu_i\}}$, meaning the utility levels and expenditures are dual (in a colloquial sense) to one another. Therefore, equilibrium utility levels can be determined from equilibrium expenditures, and vice-versa, which implies that allocations and prices can likewise be derived from one another through this duality structure.[5]

[5]A more in-depth analysis of this duality structure can be found in Blume (Blume, 2017).

Since the objective function of the primal in Theorem 6.2.1 is in general non-concave (i.e., if utilities $\boldsymbol{u}$ are not concave), strong duality need not hold; however, the dual is still guaranteed to be convex (Boyd et al., 2004). This observation suggests that even if the problem of computing Fisher equilibrium *allocations* is non-concave, the problem of computing Fisher equilibrium *prices* can still be convex. Additionally, since this convex program differs from the Eisenberg-Gale program by an additive constant, we obtain as a corollary that solutions to the Einseberg-Gale program also correspond to Fisher equilibria in *all* homothetic Fisher markets, including those in which the buyers' utility functions are non-concave.

6.3 Convex Potential Markets

An interesting property of this convex program is that its dual expresses Fisher equilibrium prices via expenditure functions, and just like the Eisenberg-Gale program's dual objective (Cheung et al., 2013; Devanur et al., 2008), the gradient of its objective $\psi(\boldsymbol{p})$ at any price $\boldsymbol{p}$ is equal to the negative excess demand in the market at those prices.

Cheung et al. (2013) showed via the Lagrangian of the Eisenberg-Gale program, i.e., without constructing the precise dual, that the subdifferential of the dual of the Eisenberg-Gale program is equal to the negative excess demand in the associated market, which implies that mirror descent equivalent to a subset of tâtonnement rules. In this section, we use a generalization of Shephard's lemma to prove that the subdifferential of the dual of our new convex program is equal to the negative excess demand in the associated market. Our proof also applies to the dual of the Eisenberg-Gale program, since the two duals differ only by a constant factor.

Shephard's lemma tells us that the rate of change in expenditure with respect to prices, evaluated at prices $\boldsymbol{p}$ and utility level ν_i, is equal to the Hicksian demand at prices $\boldsymbol{p}$ and utility level ν_i. Alternatively, the partial derivative of the expenditure function with respect to the price p_j of good j at utility level ν_i is simply the share of the total expenditure allocated to j divided by the price of j, which is exactly the Hicksian demand for j at utility level ν_i.

While Shephard's lemma is applicable to utility functions with singleton-valued Hicksian demand (i.e., strictly concave utility functions), we require a generalization of Shephard's lemma that applies to utility functions that could have set-valued Hicksian demand. An early proof of this generalized lemma was given by Tanaka (2008) in a discussion paper; a more modern perspective can be found in a recent survey by Blume (2017). For completeness, we also provide a new, simple proof of this result via Danskin's theorem (for subdifferentials) in (Danskin, 1966) Section 7.2.

Lemma 6.3.1 [Shephard's lemma, generalized for set-valued Hicksian demand (Blume, 2017; Shephard, 2015; Tanaka, 2008)]**.**
Let $e_i(\boldsymbol{p}, \nu_i)$ be the expenditure function of buyer i and $\boldsymbol{h}_i(\boldsymbol{p}, \nu_i)$ be the Hicksian demand set of buyer i. The subdifferential $\mathcal{D}_{\boldsymbol{p}} e_i(\boldsymbol{p}, \nu_i)$ is the Hicksian demand at prices $\boldsymbol{p}$ and utility level ν_i, i.e., $\mathcal{D}_{\boldsymbol{p}} e_i(\boldsymbol{p}, \nu_i) = \boldsymbol{h}_i(\boldsymbol{p}, \nu_i)$.

The next lemma plays an essential role in the proof that the subdifferential of the dual of our convex program is equal to the negative excess demand. Just as Shephard's Lemma related the expenditure function to Hicksian demand via (sub)gradients, this lemma relates the expenditure function to Marshallian demand via (sub)gradients. One way to understand this relationship is in terms of **Marshallian consumer surplus**, the area under the Marshallian demand curve, i.e., the integral of Marshallian demand with respect to prices.[6] Specifically, by applying the fundamental theorem of calculus to the left-hand side of Lemma 6.3.2, we see that the Marshallian consumer surplus equals $b_i \log \left(\frac{\partial e_i(\boldsymbol{p}, \nu_i)}{\partial \nu_i} \right)$. The key takeaway is thus that any objective function we might seek to optimize that includes a buyer's Marshallian consumer surplus is thus optimizing their Marshallian demand, so that

[6]We note that the definition of Marshallian consumer surplus for multiple goods requires great care and falls outside the scope of this paper. More information on consumer surplus can be found in Levin (2004), and Vives (1987).

optimizing this objective yields a utility-maximizing allocation for the buyer, constrained by their budget.

Lemma 6.3.2.

If buyer i's utility function u_i is continuous and homogeneous, then $\mathcal{D}_{\boldsymbol{p}}\left(b_i \log\left(\frac{\partial e_i(\boldsymbol{p},\nu_i)}{\partial \nu_i}\right)\right) = \boldsymbol{d}_i(\boldsymbol{p}, b_i)$.

Remark 6.3.1.

Lemma 6.3.2 makes the dual of our convex program easy to interpret, and thus sheds light on the dual of the Eisenberg-Gale program. Specifically, we can interpret the dual as specifying prices that minimize the distance between the sellers' surplus and the buyers' Marshallian surplus. The left hand term is simply the sellers' surplus, and by Lemma 6.3.2, the right hand term can be seen as the buyers' total Marshallian surplus.

Remark 6.3.2.

The lemmas we have proven in this section and the last provide a possible explanation as to why no primal-dual type convex program is known that solves Fisher markets when buyers have *non-homogeneous* utility functions, in which the primal describes optimal allocations while the dual describes equilibrium prices. By the homogeneity assumption, a CCH buyer can increase their utility level (resp. decrease their spending) by $c\%$ by increasing their budget (resp. decreasing their desired utility level) by $c\%$ (Lemma 7.2.1). This observation implies that the marginal expense of additional utility, i.e., "bang-per-buck", and the marginal utility of additional budget, i.e., "buck-per-bang", are constant (Lemma 6.2.1). Additionally, optimizing prices to maximize buyers' "bang-per-buck" is equivalent to optimizing prices to minimize their "buck-per-bang" (Corollary 6.2.1). Further, optimizing prices to minimize their "buck-per-bang" is equivalent to maximizing their utilities constrained by their budgets (Lemma 6.3.2). Thus, the equilibrium prices computed by the dual of our program, which optimize the buyers' buck-per-bang, simultaneously optimize their utilities constrained by their budgets. In particular, equilibrium prices can be computed without reference to equilibrium allocations (Corollary 6.2.1 + Lemma 6.3.2).

In other words, assuming homogeneity, the computation of the equilibrium allocations and prices can be isolated into separate primal and dual problems.

Next, we show that the subdifferential of the dual of our convex program is equal to the negative excess demand in the associated market.

Theorem 6.3.1.

Given any homothetic Fisher market $(\boldsymbol{u}, \boldsymbol{b})$, the subdifferential of the dual of the program in Theorem 6.2.1 at any price $\boldsymbol{p}$ is equal to the negative excess demand in $(\boldsymbol{u}, \boldsymbol{b})$ at price $\boldsymbol{p}$: i.e., $\mathcal{D}_{\boldsymbol{p}}\psi(\boldsymbol{p}) = -\mathcal{Z}(\boldsymbol{p})$.

Cheung et al. (2013) define a class of markets called **convex potential function (CPF)** markets. A market is a CPF market, if there exists a convex potential function φ such that $\mathcal{D}_{\boldsymbol{p}}\varphi(\boldsymbol{p}) = -\boldsymbol{z}(\boldsymbol{p})$. They then prove that Fisher markets are CPF markets by showing, through the Lagrangian of the Eisenberg-Gale program, that its dual is a convex potential function (Cheung et al., 2013). Likewise, Theorem 6.3.1 implies the following:

Corollary 6.3.1.

All homothetic Fisher markets are CPF markets.

Proof

A convex potential function $\phi : \mathbb{R}^m \to \mathbb{R}$ for any CCH Fisher market $(\boldsymbol{u}, \boldsymbol{b})$ is given by:

$$\psi(\boldsymbol{p}) = \sum_{j \in [m]} p_j - \sum_{i \in [n]} b_i \log\left(\frac{\partial e_i(\boldsymbol{p}, \nu_i)}{\partial \nu_i}\right) \tag{6.7}$$

6.4 Market Parameters

An important consequence of the fact that implies that mirror descent on φ over the positive ortanth is equivalent to *tâtonnement* in all homothetic Fisher markets. Using this equivalence, we can pick a particular kernel function h, and then potentially use Theorem 6.1.1 to establish convergence rates for *tâtonnement*.

Unfortunately, *tâtonnement* does not converge to equilibrium prices in all homothetic Fisher markets, e.g., linear Fisher markets (Cole and Tao, 2019), which suggests the need for additional restrictions on the class of homothetic Fisher markets. Goktas et al. (2022b) suggest the maximum absolute value of the Marshallian price demand elasticity, i.e., $c = \max_{j,k,i} \max_{(\boldsymbol{p},\boldsymbol{b}) \in \Delta_m \times \Delta_n \times [n]} \left\| \epsilon_{d_{ij}, p_k}(\boldsymbol{p}, b_i) \right\|$, as a possible market parameter to use to establish a convergence rate of $O(^{(1+c)}/_T)$. However, Cole and Fleischer's [2008] results suggest that it is unlikely that Marshallian demand elasticity could be enough, since the proof techniques used in work that makes this assumption require one to quantify the direction of the change in demand as a function of the change in the prices of the other goods, and hence only apply when one assumes WGS or WGC (Cole and Fleischer, 2008).

One of the main contributions of this paper is the observation that the maximum absolute value of the price elasticity of Hicksian demand in a homothetic Fisher market is sufficient to analyze the convergence of *tâtonnement*. To this end, in this section, we anaylyze Hicksian demand price elasticity, exposit some of its properties in homothetic markets, and argue why it is a natural parameter to consider in the analysis of *tâtonnement*.

We first note that for Leontief utilities, the Hicksian cross-price elasticity of demand is equal to 0, while for linear utilities the Hicksian cross-price elasticity of demand is, by convention, ∞.[7] For Cobb-Douglas utilities, the Hicksian cross-price elasticity of demand is strictly positive and upper bounded by 1, but it is not the same for all pairs of goods. Note that the behavior of the Hicksian cross-price elasticity of demand is radically different than that of the Marshallian cross-price elasticity of demand, for which the elasticities of linear, Cobb-Douglas, and Leontief utilities are respectively given as ∞, 0, and $-\infty$. A taxonomy of utility classes as a function of price elasticity of demand (both Marshallian and Hicksian) is shown in Figure 3.1b (Section 3.3).

[7]The limit of Hicksian price elasticity of demand as $\rho \to 1$ is not well defined, i.e., if $\rho \to 1^-$ the limit is $+\infty$, while if $\rho \to 1^+$ the limit is $-\infty$. However, for linear utilities, as the Hicksian demand for a good can only go up when the price of another good goes up, we set the elasticity of Hicksian price elasticity of demand for linear utilities to be $+\infty$, by convention. We refer the reader to Ramskov and Munksgaard (2001) for a primer on elasticity of demand.

We start our analysis with following lemma, which shows that the Hicksian price elasticity of demand is constant across all utility levels in homothetic Fisher markets. This property implies that the Hicksian demand price elasticity at one unit of utility provides sufficient information about the market's reactivity to changes in prices, even without any information about the buyers' utility levels. This information is crucial when trying to bound the changes in Hicksian demand from one iteration of *tâtonnement* to another, since buyers' utilities can change.[8]

Lemma 6.4.1.

For any Hicksian demand $\boldsymbol{h}_i$ associated with a homogeneous utility function u_i, for all $j, k \in [m], \boldsymbol{p} \in \mathbb{R}_+^m, \nu_i \in \mathbb{R}_+$, it holds that $\epsilon_{h_{ij},p_k}(\boldsymbol{p}, \nu_i) = \epsilon_{h_{ij},p_k}(\boldsymbol{p}, 1) = 1$.

With the above lemma in hand, we now explain why the Hicksian demand price elasticity[9] is a better market parameter by which to analyze the convergence of *tâtonnement* than the Marshallian demand price elasticity. Cheung et al. (2013); Cheung (2014) use the dual of the Eisenberg-Gale program as a potential to measure the progress that *tâtonnement* makes at each step, for (nested) CES and Leontief utilities. Under these functional forms, the authors are able to explain a change in the value of the buyers' indirect utilities as a function a change in prices, based on which they bound the change in the second term of the dual $\sum_{i\in[n]} b_i \log(v_i(\boldsymbol{p}, b_i)) - b_i$ from one time period to the next. Using this bound, they show that *tâtonnement* makes steady progress towards equilibrium.

However, in general homothetic Fisher markets, knowing how much the Marshallian demand for each good changes from one iteration of *tâtonnement* to another does not tell us how much the buyers' utilities change. More concretely, suppose that the Marshallian demand of a buyer i has changed by an additive vector $\Delta\boldsymbol{d}_i$ from time t to time $t+1$, then the difference in indirect utilities from one period to another is given by $u_i(\boldsymbol{d}_i^{(t+1)}) - u_i(\boldsymbol{d}_i^{(t)}) = u_i(\boldsymbol{d}_i^{(t)} + \Delta\boldsymbol{d}_i) - u_i(\boldsymbol{d}_i^{(t)})$. Without additional information about the utility functions, e.g.,

[8] We include all omitted results and proofs in Section 7.2.

[9] Going forward we refer to the Hicksian price elasticity of demand, as simply Hicksian demand elasticity, because Hicksian price elasticity of demand w.r.t. utility level is always 1.

Lipschitz continuity, it is impossible to bound this difference, because utilities can change by an unbounded amount from one period to another. Hence, even if the Marshallian price elasticity of demand and the changes in prices from one period to another were known, it would only allow us to bound the difference in demands, and not the difference in indirect utilities. To get around this difficulty, one could consider making an assumption about the boundedness of the indirect utility function's price elasticity, or the utility function's Lipschitz-continuity, but such assumptions would not be economically justified, since utility functions are merely representations of preference orderings without any inherent meaning of their own.

We can circumvent this issue by instead looking at the dual of the convex program in Theorem 6.2.1. In this dual, the indirect utility term is replaced by the expenditure function. The advantage of this formulation is that if one knows the amount by which prices change from one iteration to the next, as well as the Hicksian demand elasticity, then we can easily bound the change in spending from one period to another.

The following lemma is crucial to proving the convergence of *tâtonnement*. This lemma allows us to bound the changes in buyer spending across all time periods, thereby allowing us to obtain a global convergence rate. In particular, it shows that the change in spending between two consecutive iterations of *tâtonnement* can be bounded as a function of the prices and the Hicksian demand elasticity.

More formally, suppose that we would like to bound the percentage change in expenditure at one unit of utility from one iteration to another, i.e., $\frac{e_i(\boldsymbol{p}^{(t+1)},1)-e_i(\boldsymbol{p}^{(t)},1)}{e_i(\boldsymbol{p}^{(t)},1)}$, using a first order Taylor expansion of $e_i(\boldsymbol{p}^{(t)} + \Delta\boldsymbol{p}, 1)$ around $\boldsymbol{p}^{(t)}$. By Taylor's theorem (Graves, 1927), we have: $e_i(\boldsymbol{p}^{(t)} + \Delta\boldsymbol{p}, 1) = e_i(\boldsymbol{p}^{(t)}, 1) + \left\langle \nabla_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)}, 1), \Delta\boldsymbol{p} \right\rangle + {}^1\!/_2 \left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle$ for some $c \in (0,1)$. Re-organizing terms around, we get $e_i(\boldsymbol{p}^{(t+1)}, 1) - e_i(\boldsymbol{p}^{(t)}, 1) = \left\langle \nabla_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)}, 1), \Delta\boldsymbol{p} \right\rangle + {}^1\!/_2 \left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle$. Dividing both sides by $e_i(\boldsymbol{p}^{(t)}, 1)$

we obtain:

$$\frac{\left\langle \nabla_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)}, 1), \Delta\boldsymbol{p} \right\rangle + 1/2 \left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)}$$

We can now apply Shephard's lemma (Shephard, 2015), a corollary of the envelope theorem (Afriat, 1971; Milgrom and Segal, 2002), to the numerator, which allows us to conclude that for all buyers $i \in [n]$, $\nabla_{\boldsymbol{p}} e_i(\boldsymbol{p}, \nu_i) = \boldsymbol{h}_i(\boldsymbol{p}, \nu_i)$. Next, using the definition of the expenditure function in the denominator, we obtain the following:

$$= \frac{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \Delta\boldsymbol{p} \right\rangle}{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \boldsymbol{p}^{(t)} \right\rangle} + \frac{1/2 \left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)} \tag{6.8}$$

If the change in prices is bounded, and the Hicksian demand elasticity is known, then one can bound the first term in Equation (6.8) with ease. It remains to be seen if the second term can be bounded. The following lemma provides an affirmative answer to that question. In particular, we show that the second-order error term in the Taylor approximation above can be bounded as a function of the maximum absolute value of the Hicksian demand elasticity. We note that in the following lemma, by Lemma 7.2.9, the Marshallian demand is unique, because the Hicksian demand is a singleton for bounded elasticity of Hicksian demand.

Lemma 6.4.2.

Fix $i \in [n]$ and $t \in \mathbb{N}_+$ and let $\Delta\boldsymbol{p} = \boldsymbol{p}^{(t+1)} - \boldsymbol{p}^{(t)}$. Suppose that $\frac{|\Delta p_j|}{p_j^{(t)}} \leq \frac{1}{4}$, then for all buyers $i \in [n]$, and for some $c \in (0, 1)$, it holds that:

$$\left| \frac{b_i}{2} \frac{\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)} \right| \leq \frac{5\epsilon}{6} \sum_j \frac{(\Delta p_j)^2}{p_j^{(t)}} d_{ij}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, b_i) \ , \tag{6.9}$$

where $\epsilon \doteq \max_{\boldsymbol{p} \in \Delta_m, j,k \in [m]} \left| \epsilon_{h_{ij}, p_k}(\boldsymbol{p}, 1) \right|$.

Proof of Lemma 6.4.2

By Shephard's lemma (Shephard, 2015) (Lemma 6.3.1, Section 7.2), it holds that $\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle = \left\langle \nabla_{\boldsymbol{p}} \boldsymbol{h}_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle$.

$$\left|\frac{b_i}{2}\frac{\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle}{\left\langle \boldsymbol{h}_i^t, \boldsymbol{p}^{(t)}\right\rangle}\right|$$

$$= \left|\frac{b_i}{2}\frac{\left\langle \nabla_{\boldsymbol{p}} \boldsymbol{h}_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle}{\left\langle \boldsymbol{h}_i^t, \boldsymbol{p}^{(t)}\right\rangle}\right| \qquad \text{(Shephard's Lemma)}$$

$$\leq \frac{b_i}{2}\frac{\sum_{j,k}\left|\Delta p_j\right|\left|\mathcal{D}_{p_k} h_{ij}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\right|\left|\Delta p_k\right|}{\left\langle \boldsymbol{h}_i^t, \boldsymbol{p}^{(t)}\right\rangle}$$

$$= \frac{b_i}{2}\frac{\sum_{j,k}\left|\Delta p_j\right|\sqrt{\left|\mathcal{D}_{p_k} h_{ij}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\right|}\sqrt{\left|\mathcal{D}_{p_k} h_{ij}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\right|}\left|\Delta p_k\right|}{\left\langle \boldsymbol{h}_i^t, \boldsymbol{p}^{(t)}\right\rangle}$$

$$= \frac{b_i}{2}\frac{\sum_{j,k}\left|\Delta p_j\right|\sqrt{\left|\mathcal{D}_{p_j} h_{ik}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\right|}\sqrt{\left|\mathcal{D}_{p_k} h_{ij}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\right|}\left|\Delta p_k\right|}{\left\langle \boldsymbol{h}_i^t, \boldsymbol{p}^{(t)}\right\rangle} \tag{6.10}$$

where the last was obtained from the symmetry of $\nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}, \nu_i) = \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}, \nu_i)^T$ for all $i \in [n], \boldsymbol{p} \in \mathbb{R}^m_+, \nu_i \in \mathbb{R}_+$ (Mas-Colell et al., 1995), which combined with Shephard's lemma gives us $\nabla_{\boldsymbol{p}} \boldsymbol{h}_i(\boldsymbol{p}, \nu_i) = \nabla_{\boldsymbol{p}} \boldsymbol{h}_i(\boldsymbol{p}, \nu_i)^T$, i.e., for all $j, k \in [m]$, $\mathcal{D}_{p_j} h_{ik}(\boldsymbol{p}, \nu_i) = \mathcal{D}_{p_k} h_{ij}(\boldsymbol{p}, \nu_i)$.

Define the Hicksian demand elasticity of buyer i for good j w.r.t. the price of good k as $\epsilon_{h_{ij}, p_k}(\boldsymbol{p}, \nu_i) = \mathcal{D}_{p_k} h_{ij}(\boldsymbol{p}, \nu_i)\frac{p_k}{h_{ij}(\boldsymbol{p}, \nu_i)}$. Since utility functions are homogeneous, by Lemma 6.4.1 we have for all $\nu_i \in \mathbb{R}_+$, $\epsilon_{h_{ij}, p_k}(\boldsymbol{p}, \nu_i) = \mathcal{D}_{p_k} h_{ij}(\boldsymbol{p}, \nu_i)\frac{p_k}{h_{ij}(\boldsymbol{p}, \nu_i)} = \mathcal{D}_{p_k} h_{ij}(\boldsymbol{p}, 1)\frac{p_k}{h_{ij}(\boldsymbol{p}, 1)}$. Re-organizing expressions, we get $\mathcal{D}_{p_k} h_{ij}(\boldsymbol{p}, 1) = \epsilon_{h_{ij}, p_k}(\boldsymbol{p}, 1)\frac{h_{ij}(\boldsymbol{p}, 1)}{p_k}$. Going back to Equation (6.10), we get:

$$= \frac{b_i}{2}\frac{\sum_{j,k}\left|\Delta p_j\right|\sqrt{\left|\epsilon_{h_{ik}, p_j}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\frac{h_{ik}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)}{p_j^{(t)} + c\Delta p_j}\right|}\sqrt{\left|\epsilon_{h_{ij}, p_k}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\frac{h_{ij}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)}{p_k^{(t)} + c\Delta p_k}\right|}\left|\Delta p_k\right|}{\left\langle \boldsymbol{h}_i^t, \boldsymbol{p}^{(t)}\right\rangle}$$

$$= \frac{b_i}{2}\frac{\sum_{j,k}\left|\Delta p_j\right|\sqrt{\left|\epsilon_{h_{ik}, p_j}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\right|\frac{h_{ik}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)}{p_j^{(t)} + c\Delta p_j}}\sqrt{\left|\epsilon_{h_{ij}, p_k}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\right|\frac{h_{ij}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)}{p_k^{(t)} + c\Delta p_k}}\left|\Delta p_k\right|}{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \boldsymbol{p}^{(t)}\right\rangle}$$

Letting $\epsilon = \max_{\boldsymbol{p} \in \mathbb{R}^m_+, \nu_i \in \mathbb{R}_+, j,k \in [m]}\left|\epsilon_{h_{ij}, p_k}(\boldsymbol{p}, \nu_i)\right|$. Note that since utility functions are homogeneous, by Lemma 6.4.1 we have $\epsilon = \max_{\boldsymbol{p} \in \mathbb{R}^m_+, \nu_i \in \mathbb{R}_+, j,k \in [m]}\left|\epsilon_{h_{ij}, p_k}(\boldsymbol{p}, \nu_i)\right| = \max_{\boldsymbol{p} \in \mathbb{R}^m_+, j,k \in [m]}\left|\epsilon_{h_{ij}, p_k}(\boldsymbol{p}, 1)\right|$, which gives us:

$$\leq \frac{\epsilon b_i}{2}\frac{\sum_{j,k}\left|\Delta p_j\right|\sqrt{\frac{h_{ik}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)}{p_j^{(t)} + c\Delta p_j}}\sqrt{\frac{h_{ij}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)}{p_k^{(t)} + c\Delta p_k}}\left|\Delta p_k\right|}{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \boldsymbol{p}^{(t)}\right\rangle}$$

Since for all $j \in [m]$, $\frac{|\Delta p_j|}{p_j} \leq \frac{1}{4}$, we have that for all $j \in [m]$ and for all $c \in [0,1]$, $p_j^{(t)} + c\Delta p_j \geq \frac{3}{4} p_j^{(t)}$, which gives:

$$
\begin{aligned}
&\leq \frac{\epsilon b_i}{2} \frac{\sum_{j,k} \left|\Delta p_j\right| \sqrt{\frac{h_{ik}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p},1)}{\frac{3}{4}p_j^{(t)}}} \sqrt{\frac{h_{ij}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p},1)}{\frac{3}{4}p_k^{(t)}}} \left|\Delta p_k\right|}{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)},1), \boldsymbol{p}^{(t)} \right\rangle} \\
&= \frac{2\epsilon b_i}{3} \frac{\sum_{j,k} \left|\Delta p_j\right| \sqrt{\frac{h_{ij}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p},1)}{p_j^{(t)}}} \sqrt{\frac{h_{ik}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p},1)}{p_k^{(t)}}} \left|\Delta p_k\right|}{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)},1), \boldsymbol{p}^{(t)} \right\rangle} \\
&= \frac{2\epsilon b_i}{3} \frac{\sum_{j,k\in[m]} \sqrt{\frac{\left|\Delta p_j\right|^2}{p_j^{(t)}} h_{ij}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p},1) h_{ik}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p},1) \frac{\left|\Delta p_k\right|^2}{p_k^{(t)}}}}{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)},1), \boldsymbol{p}^{(t)} \right\rangle}
\end{aligned}
$$

Applying the AM-GM inequality, i.e., for all $x, y \in \mathbb{R}_+$, $\frac{x+y}{2} \geq \sqrt{xy}$, to the sum inside the numerator above, we obtain:

$$
\begin{aligned}
&\leq \frac{2\epsilon b_i}{3} \frac{\sum_{j,k\in[m]} 1/2 \left(\frac{\left|\Delta p_j\right|^2}{p_j^{(t)}} h_{ij}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p},1) + h_{ik}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p},1) \frac{\left|\Delta p_k\right|^2}{p_k^{(t)}} \right)}{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)},1), \boldsymbol{p}^{(t)} \right\rangle} \\
&\leq \frac{2\epsilon b_i}{3} \frac{\sum_j \frac{(\Delta p_j)^2}{p_j^{(t)}} h_{ij}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p},1)}{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)},1), \boldsymbol{p}^{(t)} \right\rangle}
\end{aligned}
$$

Since for all $j \in [m]$, $\frac{|\Delta p_j|}{p_j^{(t)}} \leq \frac{1}{4}$, we have for all $c \in [0,1]$ that $\frac{4}{5} \sum_j h_{ij}(\boldsymbol{p}^{(t)},1)(p_j^{(t)} + c\Delta p_j) \leq \sum_j h_{ij}(\boldsymbol{p}^{(t)},1) p_j^{(t)}$:

$$
\begin{aligned}
&\leq \frac{2\epsilon b_i}{3} \frac{\sum_j \frac{(\Delta p_j)^2}{p_j^{(t)}} h_{ij}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p},1)}{\frac{4}{5}\sum_j h_{ij}(\boldsymbol{p}^{(t)},1)(p_j^{(t)} + c\Delta p_j)} \\
&= \frac{5\epsilon b_i}{6} \frac{\sum_j \frac{(\Delta p_j)^2}{p_j^{(t)}} h_{ij}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p},1)}{\sum_j h_{ij}(\boldsymbol{p}^{(t)},1)(p_j^{(t)} + c\Delta p_j)} \\
&\leq \frac{5\epsilon b_i}{6} \frac{\sum_j \frac{(\Delta p_j)^2}{p_j^{(t)}} h_{ij}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p},1)}{\sum_j h_{ij}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p},1)(p_j^{(t)} + c\Delta p_j)} && \text{(Corollary 7.2.1, Section 7.2)} \\
&= \frac{5\epsilon}{6} \sum_j \frac{(\Delta p_j)^2}{p_j^{(t)}} d_{ij}(\boldsymbol{p}^{(t)}+c\Delta \boldsymbol{p}, b_i) && \text{(Lemma 7.2.9, Section 7.2)}
\end{aligned}
$$

Because we can bound the change in the expenditure function from one iteration of *tâtonnement* to the next, the Hicksian price elasticity of demand is a better tool with which to analyze the convergence of *tâtonnement* than Marshallian price elasticity of demand. Additionally, as shown previously by Cheung et al. (2013) (Lemma 7.2.4, Section 7.2), we can further upper bound the price terms in Lemma 6.4.2 by the KL divergence between the two prices In light of Theorem 6.1.1, this result suggests that running mirror descent with KL divergence as the Bregman divergence on the dual of the convex program in Theorem 6.2.1 could result in a *tâtonnement* update rule that converges to a Walrasian equilibrium.

6.5 Convergence Bounds for Entropic Tâtonnement

In this section, we analyze the rate of convergence of **entropic tâtonnement**, which corresponds to the *tâtonnement* process given by mirror descent with weighted entropy as the kernel function, i.e., entropic descent. This particular update rule reduces to Equations (6.4) to (6.5), and has been the focus of previous work (Cheung et al., 2013). We provide a sketch of the proof used to obtain our convergence rate in this section. The omitted lemmas and proofs can be found in Appendix 7.2.

At a high level, our proof follows Cheung et al.'s [2013] proof technique for Leontief Fisher markets (Cheung et al., 2013), although we encounter different lower-level technical challenges in generalizing to homothetic markets. This proof technique works as follows. First, we prove that under certain assumptions, the condition required by Theorem 6.1.1 holds when f is the convex potential function for homothetic Fisher markets, i.e., $f = \psi$. For these assumptions to be valid, we need to set γ to be greater than a quadratic function of the maximum absolute value of the price elasticity of the Hicksian demand and the maximum Marshallian demand, for all goods throughout the *tâtonnment* process. Further, since γ needs to be set at the outset, we need to upper bound γ. To do so, we derive a bound on the maximum demand for any good during *tâtonnement* in all homothetic Fisher markets, which in turn allows us to derive an upper bound on γ. Finally, we use Theorem 6.1.1 to obtain the convergence rate of $O(1+\epsilon^2/t)$.

The following lemma derives the conditions under which the antecedent of Theorem 6.1.1 holds for entropic *tâtonnement*.

Lemma 6.5.1.

Consider a homothetic Fisher market $(\boldsymbol{u}, \boldsymbol{b})$ and let $\epsilon = \max_{\boldsymbol{p} \in \Delta_m, j,k \in [m]} \left| \epsilon_{h_{ij}, p_k}(\boldsymbol{p}, 1) \right|$. Then, the following holds for entropic *tâtonnement* when run on $(\boldsymbol{u}, \boldsymbol{b})$: for all $t \in \mathbb{N}$,

$$\psi(\boldsymbol{p}^{(t+1)}) \leq \ell_\psi(\boldsymbol{p}^{(t+1)}, \boldsymbol{p}^{(t)}) + \gamma \mathrm{div}_{\mathrm{KL}}(\boldsymbol{p}^{(t+1)}, \boldsymbol{p}^{(t)}) \ ,$$

where $\gamma = \left(1 + \max_{j \in [m]} \left\{ \sum_{i \in [n]} \max_{t \in \mathbb{N}_+} v_i(\boldsymbol{p}^{(t)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) \right\} \right) \left(6 + \frac{85\epsilon}{12} + \frac{25\epsilon^2}{72}\right)$.

Proof of Lemma 6.5.1

$$
\begin{aligned}
&\psi(\boldsymbol{p}^{(t+1)}) - \ell_\psi(\boldsymbol{p}^{(t+1)}, \boldsymbol{p}^{(t)}) \\
&= \psi(\boldsymbol{p}^{(t+1)}) - \psi(\boldsymbol{p}^{(t)}) + \boldsymbol{z}(\boldsymbol{p}^{(t)}) \cdot \left(\boldsymbol{p}^{(t+1)} - \boldsymbol{p}^{(t)}\right) \\
&= \sum_{j \in [m]} \left(p_j^{(t)} + \Delta p_j\right) - \sum_{i \in [n]} b_i \log\left(e_i(\boldsymbol{p}^{(t+1)}, 1)\right) - \sum_{j \in [m]} p_j^{(t)} + \sum_{i \in [n]} b_i \log\left(e_i(\boldsymbol{p}^{(t)}, 1)\right) + \sum_{j \in [m]} z_j(\boldsymbol{p}^{(t)}) \Delta p_j \\
&= \sum_{j \in [m]} \left(p_j^{(t)} + \Delta p_j\right) - \sum_{i \in [n]} b_i \log\left(e_i(\boldsymbol{p}^{(t+1)}, 1)\right) - \sum_{j \in [m]} p_j^{(t)} + \sum_{i \in [n]} b_i \log\left(e_i(\boldsymbol{p}^{(t)}, 1)\right) + \sum_{j \in [m]} (q_j^{(t)} - 1) \Delta p_j \\
&= \sum_{j \in [m]} \Delta p_j q_j^{(t)} - \sum_{i \in [n]} b_i \log\left(e_i(\boldsymbol{p}^{(t+1)}, 1)\right) + \sum_{i \in [n]} b_i \log\left(e_i(\boldsymbol{p}^{(t)}, 1)\right) \\
&= \left\langle \Delta \boldsymbol{p}, \boldsymbol{q}^{(t)} \right\rangle + \sum_{i \in [n]} b_i \log\left(\frac{e_i(\boldsymbol{p}^{(t)}, 1)}{e_i(\boldsymbol{p}^{(t+1)}, 1)}\right) \\
&= \left\langle \Delta \boldsymbol{p}, \boldsymbol{q}^{(t)} \right\rangle + \sum_{i \in [n]} b_i \log\left(\frac{e_i(\boldsymbol{p}^{(t)}, 1)}{e_i(\boldsymbol{p}^{(t)}, 1) + e_i(\boldsymbol{p}^{(t+1)}, 1) - e_i(\boldsymbol{p}^{(t)}, 1)}\right) \\
&= \left\langle \Delta \boldsymbol{p}, \boldsymbol{q}^{(t)} \right\rangle + \sum_{i \in [n]} b_i \log\left(1 - \frac{e_i(\boldsymbol{p}^{(t+1)}, 1) - e_i(\boldsymbol{p}^{(t)}, 1)}{e_i(\boldsymbol{p}^{(t)}, 1)} \left(1 + \frac{e_i(\boldsymbol{p}^{(t+1)}, 1) - e_i(\boldsymbol{p}^{(t)}, 1)}{e_i(\boldsymbol{p}^{(t)}, 1)}\right)^{-1}\right)
\end{aligned}
$$

where the last line is obtained by simply noting that $\forall a, b \in \mathbb{R}$, $\frac{a}{a+b} = 1 - \frac{b}{a}(1+\frac{b}{a})^{-1}$. Using Lemma 7.2.12, we then obtain:

$$\psi(\boldsymbol{p}^{(t+1)}) - \ell_\psi(\boldsymbol{p}^{(t+1)}, \boldsymbol{p}^{(t)})$$

$$\leq \left\langle \Delta\boldsymbol{p}, \boldsymbol{q}^{(t)}\right\rangle + \sum_{i\in[n]}\left[\left(\frac{4}{3}+\frac{20\epsilon}{27}\right)\sum_{l\in[m]}\frac{d_{il}^{(t)}}{p_l^{(t)}}(\Delta p_l)^2 + \left(\frac{5\epsilon}{6}+\frac{25\epsilon^2}{324}\right)\sum_{l\in[m]}\frac{d_{il}(\boldsymbol{p}^{(t)}+c\Delta\boldsymbol{p}, b_i)}{p_l^{(t)}}(\Delta p_l)^2 - \left\langle \boldsymbol{d}_i^{(t)}, \Delta\boldsymbol{p}\right\rangle\right]$$

$$= \left\langle \Delta\boldsymbol{p}, \boldsymbol{q}^{(t)}\right\rangle + \left(\frac{4}{3}+\frac{20\epsilon}{27}\right)\sum_{l\in[m]}\frac{q_l^{(t)}}{p_l^{(t)}}(\Delta p_l)^2 + \left(\frac{5\epsilon}{6}+\frac{25\epsilon^2}{324}\right)\sum_{l\in[m]}\frac{q_l(\boldsymbol{p}^{(t)}+c\Delta\boldsymbol{p}, \boldsymbol{b})}{p_l^{(t)}}(\Delta p_l)^2 - \left\langle \boldsymbol{q}^{(t)}, \Delta\boldsymbol{p}\right\rangle$$

$$= \left(\frac{4}{3}+\frac{20\epsilon}{27}\right)\sum_{l\in[m]}\frac{q_l^{(t)}}{p_l^{(t)}}(\Delta p_l)^2 + \left(\frac{5\epsilon}{6}+\frac{25\epsilon^2}{324}\right)\sum_{l\in[m]}\frac{q_l(\boldsymbol{p}^{(t)}+c\Delta\boldsymbol{p}, \boldsymbol{b})}{p_l^{(t)}}(\Delta p_l)^2$$

$$\leq \left(\frac{4}{3}+\frac{20\epsilon}{27}\right)\sum_{l\in[m]}q_l^{(t)}\left(\frac{9}{2}\right)\mathrm{div}_{\mathrm{KL}}(p_j^{(t)}+\Delta p_j, p_j^{(t)})$$

$$+\left(\frac{5\epsilon}{6}+\frac{25\epsilon^2}{324}\right)\sum_{l\in[m]}q_l(\boldsymbol{p}^{(t)}+c\Delta\boldsymbol{p}, \boldsymbol{b})\left(\frac{9}{2}\right)\mathrm{div}_{\mathrm{KL}}(p_j^{(t)}+\Delta p_j, p_j^{(t)})$$

where the last line follows from Lemma 7.2.4. Continuing,

$$= \left(6+\frac{10\epsilon}{3}\right)\sum_{l\in[m]}q_l^{(t)}\mathrm{div}_{\mathrm{KL}}(p_l^{(t)}+\Delta p_l, p_l^{(t)}) \tag{6.11}$$

$$+\left(\frac{15\epsilon}{4}+\frac{25\epsilon^2}{72}\right)\sum_{l\in[m]}q_l(\boldsymbol{p}^{(t)}+c\Delta\boldsymbol{p}, \boldsymbol{b})\mathrm{div}_{\mathrm{KL}}(p_l^{(t)}+\Delta p_l, p_l^{(t)})$$

$$\leq \max_{\substack{j\in[m],\\ t\in\mathbb{N}_+}}\left\{q_j^{(t)}\right\}\left(6+\frac{10\epsilon}{3}\right)\sum_{l\in[m]}\mathrm{div}_{\mathrm{KL}}(p_l^{(t)}+\Delta p_l, p_l^{(t)}) +$$

$$\max_{\substack{j\in[m],\\ t\in\mathbb{N}_+,\\ c\in[0,1]}}\left\{q_j(\boldsymbol{p}^{(t)}+c\Delta\boldsymbol{p})\right\}\left(\frac{15\epsilon}{4}+\frac{25\epsilon^2}{72}\right)\sum_{l\in[m]}\mathrm{div}_{\mathrm{KL}}(p_l^{(t)}+\Delta p_l, p_l^{(t)})$$

$$\leq \max_{j\in[m], t\in\mathbb{N}_+, c\in[0,1]}\left\{q_j(\boldsymbol{p}^{(t)}+c\Delta\boldsymbol{p})\right\}\left(6+\frac{85\epsilon}{12}+\frac{25\epsilon^2}{72}\right)\mathrm{div}_{\mathrm{KL}}(\boldsymbol{p}^{(t)}+\Delta\boldsymbol{p}, \boldsymbol{p}^{(t)}) \tag{6.12}$$

By Lemma 7.2.9, we can rewrite the aggregate demand for all goods $j \in [m]$, as follows:

$$q_j(\boldsymbol{p}^{(t)}+c\Delta\boldsymbol{p}) = \sum_{i\in[n]} d_{ij}((1-c)\boldsymbol{p}^{(t)}+c\boldsymbol{p}^{(t+1)}, b_i)$$

$$= \sum_{i\in[n]}\frac{h_{ij}((1-c)\boldsymbol{p}^{(t)}+c\boldsymbol{p}^{(t+1)}, 1)b_i}{e_i((1-c)\boldsymbol{p}^{(t)}+c\boldsymbol{p}^{(t+1)}, 1)}$$

Now, by Danskin's maximum theorem (Danskin, 1966), we know that the expenditure function is concave in prices, that is, for all $c \in [0, 1]$, we have $(1-c)e_i(\boldsymbol{p}^{(t)}, 1) +$

$ce_i(\boldsymbol{p}^{(t+1)}, 1) \leq e_i((1-c)\boldsymbol{p}^{(t)} + c\boldsymbol{p}^{(t+1)}, 1)$. Hence, continuing we have for all $j \in [m]$:

$$q_j(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}) = \sum_{i\in[n]} \frac{h_{ij}((1-c)\boldsymbol{p}^{(t)} + c\boldsymbol{p}^{(t+1)}, 1)b_i}{e_i((1-c)\boldsymbol{p}^{(t)} + c\boldsymbol{p}^{(t+1)}, 1)}$$
$$\leq \sum_{i\in[n]} \frac{h_{ij}((1-c)\boldsymbol{p}^{(t)} + c\boldsymbol{p}^{(t+1)}, 1)b_i}{(1-c)e_i(\boldsymbol{p}^{(t)}, 1) + ce_i(\boldsymbol{p}^{(t+1)}, 1)}$$

Further, by Lemma 5 of Goktas et al. (2022b), since the expenditure function is homogeneous of degree 0 in prices, notice that we have for all $j \in [m]$, $\max_{\boldsymbol{p}\in\mathbb{R}^m_+/\{\boldsymbol{0}\}} h_{ij}(\boldsymbol{p}, 1) = \max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1)$. Note that $\max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1)$ is well-defined since Δ_m is compact, $h_{ij}(\boldsymbol{p}, 1)$ exists for all $\boldsymbol{p} \in \mathbb{R}^m_+$, and is by Berge's maximum theorem (Berge, 1997) continuous in homothetic Fisher markets. Since by the entropic tâtonnement update rule for all time-steps $t \in \mathbb{N}_+$, and goods $j \in [m]$, $p_j^{(t)} > 0$, we then have $h_{ij}((1-c)\boldsymbol{p}^{(t)} + c\boldsymbol{p}^{(t+1)}, 1) \leq \max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1)$. Hence, continuing, we have for all $j \in [m]$:

$$q_j(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}) \leq \sum_{i\in[n]} \frac{\max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1)b_i}{(1-c)e_i(\boldsymbol{p}^{(t)}, 1) + ce_i(\boldsymbol{p}^{(t+1)}, 1)}$$

Taking a maximum over $c \in [0, 1]$ and $t \in \mathbb{N}_+$, and $j \in [m]$, we have for all goods $j \in [m]$:

$$\max_{j\in[m],t\in\mathbb{N}_+,c\in[0,1]} q_j(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}) \leq \max_{j\in[m],t\in\mathbb{N}_+,c\in[0,1]} \sum_{i\in[n]} \frac{\max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1)b_i}{(1-c)e_i(\boldsymbol{p}^{(t)}, 1) + ce_i(\boldsymbol{p}^{(t+1)}, 1)}$$
$$\leq \max_{j\in[m]} \sum_{i\in[n]} \frac{\max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1)b_i}{\min_{t\in\mathbb{N}_+,c\in[0,1]}\{(1-c)e_i(\boldsymbol{p}^{(t)}, 1) + ce_i(\boldsymbol{p}^{(t+1)}, 1)\}}$$
$$= \max_{j\in[m]} \sum_{i\in[n]} \frac{\max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1)b_i}{\min_{t\in\mathbb{N}_+}\{\min\{e_i(\boldsymbol{p}^{(t)}, 1), e_i(\boldsymbol{p}^{(t+1)}, 1)\}\}}$$
$$= \max_{j\in[m]} \sum_{i\in[n]} \frac{\max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1)b_i}{\min_{t\in\mathbb{N}_+} e_i(\boldsymbol{p}^{(t)}, 1)}$$
$$= \max_{j\in[m]} \sum_{i\in[n]} \max_{t\in\mathbb{N}_+} \frac{\max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1)b_i}{e_i(\boldsymbol{p}^{(t)}, 1)}$$
$$= \max_{j\in[m]} \sum_{i\in[n]} \max_{t\in\mathbb{N}_+} v_i(\boldsymbol{p}^{(t)}, b_i) \max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1)$$

where the last line follows from Corollary 1, Appendix A of Goktas et al. (2022b). Plugging the above bound into Equation (6.12), we then obtain the following bound which implies the result:

$$\psi(\boldsymbol{p}^{(t+1)}) - \ell_\psi(\boldsymbol{p}^{(t+1)}, \boldsymbol{p}^{(t)})$$
$$\leq \max_{j \in [m]} \left\{ \sum_{i \in [n]} \max_{t \in \mathbb{N}_+} v_i(\boldsymbol{p}^{(t)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) \right\} \left(6 + \frac{85\epsilon}{12} + \frac{25\epsilon^2}{72}\right) \mathrm{div}_{\mathrm{KL}}(\boldsymbol{p}^{(t)} + \Delta\boldsymbol{p}, \boldsymbol{p}^{(t)})$$

For the above lemma to be applied in conjuction with Theorem 6.1.1, we have to ensure that the quantity $\max_{j \in [m], t \in \mathbb{N}_+} \left\{ \sum_{i \in [n]} v_i(\boldsymbol{p}^{(t)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) \right\}$ is bounded throughout entropic *tâtonnement* for homothetic Fisher markets. To understand the relevance of this bound, we note that this quantity is an upper bound to the aggregate demand, that is:

$$\begin{aligned}
q_j^{(t)} &= \sum_{i \in [n]} d_{ij}^{(t)} \\
&= \sum_{i \in [n]} h_{ij}(\boldsymbol{p}^{(t)}, v_i(\boldsymbol{p}^{(t)}, b_i)) \\
&= \sum_{i \in [n]} v_i(\boldsymbol{p}^{(t)}, b_i) h_{ij}(\boldsymbol{p}^{(t)}, 1) && \text{(Lemma 5 of Goktas et al. (2022b))} \\
&\leq \sum_{i \in [n]} \min_{t \in \mathbb{N}_+} v_i(\boldsymbol{p}^{(t)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1)
\end{aligned}$$

As such, proving an upper bound to it implies the excess demand is bounded throughout entropic *tâtonnement*, which in turn implies Lipschitz-smoothness (and hence Bregman-smoothness for any choice of strongly convex kernel function) of the dual of our convex program over all trajectories of entropic *tâtonnement*. The following lemma establishes such a bound and shows that it depends on the initial choice of price $\boldsymbol{p}^{(0)}$, and the maximum possible Hicksian demand to obtain one unit of utility.

Lemma 6.5.2 [Bounded Indirect Utility for Homothetic Fisher Markets]**.**

If entropic *tâtonnement* is run on a homothetic Fisher market $(\boldsymbol{u}, \boldsymbol{b})$, then, for all $t \in \mathbb{N}_+$, the following bound holds:

$$v_i(\boldsymbol{p}^{(t)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) \leq v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) + 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)^2}{h_{ik}(\boldsymbol{p}, 1)^2} \right\}$$

Proof of Lemma 6.5.2

Fix a buyer $i \in [n]$ and good $j \in [m]$. First, note that since by Lemma 5 of Goktas et al. (2022b), since the expenditure function is homogeneous of degree 0 in prices, we have for all $j \in [m]$, $\max_{\boldsymbol{p} \in \mathbb{R}^m_+ / \{\boldsymbol{0}\}} h_{ij}(\boldsymbol{p}, 1) = \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1)$. In addition, note that $\max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1)$ is well-defined since Δ_m is compact, $h_{ij}(\boldsymbol{p}, 1)$ exists for all $\boldsymbol{p} \in \mathbb{R}^m_+$, and is by Berge's maximum theorem (Berge, 1997) continuous in homothetic Fisher markets. Further, by the entropic tâtonnement update rule for all time-steps $t \in \mathbb{N}_+$, and goods $j \in [m]$, $p_j^{(t)} > 0$, we then have $h_{ij}(\boldsymbol{p}^{(t)}, 1) \leq \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1)$. We now proceed to prove the claim of the lemma by induction on t.

Base case: $t = 0$. Since $\max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)^2}{h_{ik}(\boldsymbol{p}, 1)^2} \right\} \geq 0$, by definition, we have

$$v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) \leq v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) + 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)^2}{h_{ik}(\boldsymbol{p}, 1)^2} \right\} .$$

Inductive hypothesis. Suppose that for any $t \in \mathbb{N}$, we have:

$$v_i(\boldsymbol{p}^{(t)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) \leq v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) + 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)^2}{h_{ik}(\boldsymbol{p}, 1)^2} \right\}$$

Inductive step. We will show that the inductive hypothesis holds for $t + 1$. We proceed with a proof by cases.

Case 1: $d_{ij}^{(t)} \geq \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)}{h_{ik}(\boldsymbol{p}, 1)} \right\}$.

For all $k \in [m]$, we have:

$$d_{ik}^{(t)} = \frac{h_{ik}^{(t)}}{h_{ij}^{(t)}} d_{ij}^{(t)} \qquad \text{(Lemma 7.2.9, Section 7.2)}$$
$$\geq \frac{h_{ik}^{(t)}}{h_{ij}^{(t)}} \max_{\substack{\boldsymbol{p},\boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p},1)>0}} \left\{ \frac{h_{ij}(\boldsymbol{q},1)}{h_{ik}(\boldsymbol{p},1)} \right\}$$
$$\geq \frac{h_{ik}^{(t)}}{h_{ij}^{(t)}} \frac{h_{ij}^{(t)}}{h_{ik}^{(t)}} = 1$$

where the penultimate line follows from the case hypothesis.

The above means that the price of all goods will increase in the next time period, i.e., $\forall k \in [m], p_k^{(t+1)} \geq p_k^{(t)}$ which implies that $e_i(\boldsymbol{p}^{(t+1)}, 1) \geq e_i(\boldsymbol{p}^{(t)}, 1) \geq 0$. In addition, note that the expenditure is positive since prices reach 0 only asymptotically under entropic *tâtonnement*. Which gives us:

$$\frac{b_i}{e_i(\boldsymbol{p}^{(t+1)}, 1)} \leq \frac{b_i}{e_i(\boldsymbol{p}^{(t)}, 1)}$$
$$v_i(\boldsymbol{p}^{(t+1)}, b_i) \leq v_i(\boldsymbol{p}^{(t)}, b_i) \qquad \text{(Corollary 1 of Goktas et al. (2022b))}$$

Multiplying both sides by $\max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1)$, we have for all $j \in [m]$:

$$v_i(\boldsymbol{p}^{(t+1)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) \leq v_i(\boldsymbol{p}^{(t)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1)$$
$$= v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) + 2 \max_{\substack{\boldsymbol{p},\boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p},1)>0}} \left\{ \frac{h_{ij}(\boldsymbol{q},1)^2}{h_{ik}(\boldsymbol{p},1)^2} \right\}$$

where the last line follows by the induction hypothesis.

Case 2: $d_{ij}^{(t)} < \max_{\substack{\boldsymbol{p},\boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p},1)>0}} \left\{ \frac{h_{ij}(\boldsymbol{q},1)}{h_{ik}(\boldsymbol{p},1)} \right\}$.

For all $k \in [m]$, we have:

$$d_{ik}^{(t)} = \frac{h_{ik}^{(t)}}{h_{ij}^{(t)}} d_{ij}^{(t)}$$
$$\leq \frac{h_{ik}^{(t)}}{h_{ij}^{(t)}} \max_{\substack{\boldsymbol{p},\boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p},1)>0}} \left\{ \frac{h_{ij}(\boldsymbol{q},1)}{h_{ik}(\boldsymbol{p},1)} \right\}$$
$$= \frac{h_{ik}^{(t)}}{h_{ij}^{(t)}} \frac{h_{ij}^{(t)}}{h_{ik}^{(t)}} = 1$$

where the penultimate line follows from the case hypothesis.

The above means that prices of all goods will decrease in the next time period. Now, note that regardless of the aggregate demand $\boldsymbol{q}^{(t)}$ at time $t \in \mathbb{N}$, prices can decrease at most by a factor of $e^{-\frac{1}{5}} \geq 1/2$, that is, for all $j \in [m]$

$$\begin{aligned}
p_j^{(t+1)} &= p_j^{(t)} \exp\left\{\frac{z_j(\boldsymbol{p}^{(t)})}{\gamma}\right\} \\
&= p_j^{(t)} \exp\left\{\frac{q_j^{(t)} - 1}{\gamma}\right\} \\
&\geq p_j^{(t)} \exp\left\{\frac{-1}{\gamma}\right\} \\
&\geq p_j^{(t)} \exp\left\{\frac{-1}{5 \max\limits_{\substack{t \in \mathbb{N} \\ j \in [m]}} \{1, q_j^{(t)}\}}\right\} \\
&\geq p_j^{(t)} \exp\left\{\frac{-1}{5}\right\} \\
&\geq p_j^{(t)} e^{-\frac{1}{5}} \geq \frac{1}{2} p_j^{(t)}
\end{aligned}$$

Now, notice that we have $e_i(\boldsymbol{p}^{(t+1)}, 1) \geq e_i(\frac{1}{2}\boldsymbol{p}^{(t)}, 1) = \frac{1}{2} e_i(\boldsymbol{p}^{(t)}, 1) \geq 0$, since the expenditure of the buyer decreases the most when the prices of all goods decrease simultaneously and the expenditure function is homogeneous of degree 1 in prices. In addition, note that the expenditure is positive since prices reach 0 only asymptotically under entropic *tâtonnement*. Hence, we have:

$$\begin{aligned}
&\frac{b_i}{e_i(\boldsymbol{p}^{(t+1)}, 1)} \leq 2 \frac{b_i}{e_i(\boldsymbol{p}^{(t)}, 1)} \\
&v_i(\boldsymbol{p}^{(t+1)}, b_i) \leq 2 v_i(\boldsymbol{p}^{(t)}, b_i) && \text{(Corollary 1 of Goktas et al. (2022b))}
\end{aligned}$$

Multiplying both sides by $h_{ij}^{(t)}$, and applying Lemma 7.2.9, we have for all $j \in [m]$:

$$\begin{aligned}
v_i(\boldsymbol{p}^{(t+1)}, b_i) h_{ij}^{(t)} &\leq 2 d_{ij}^{(t)} \\
v_i(\boldsymbol{p}^{(t+1)}) h_{ij}^{(t)} &\leq 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m] : h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{\frac{h_{ij}(\boldsymbol{q}, 1)}{h_{ik}(\boldsymbol{p}, 1)}\right\} && \text{(Case hypothesis)}
\end{aligned}$$

Now, taking a minimum over all $j \in [m]$ s.t. $h_{ij}^{(t)} > 0$, we have

$$v_i(\boldsymbol{p}^{(t+1)}) \min_{k \in [m]: h_{ik}^{(t)} > 0} h_{ik}^{(t)} \leq 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)}{h_{ik}(\boldsymbol{p}, 1)} \right\}$$

$$v_i(\boldsymbol{p}^{(t+1)}) \min_{\substack{\boldsymbol{p} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} h_{ik}(\boldsymbol{p}, 1) \leq 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)}{h_{ik}(\boldsymbol{p}, 1)} \right\}$$

$$v_i(\boldsymbol{p}^{(t+1)}) \leq \frac{2}{\min_{\substack{\boldsymbol{p} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} h_{ik}(\boldsymbol{p}, 1)} \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)}{h_{ik}(\boldsymbol{p}, 1)} \right\}$$

Finally, multiplying both sides by $\max_{\boldsymbol{q} \in \Delta_m} h_{ij}(\boldsymbol{q}, 1)$, we have:

$$\begin{aligned}
& v_i(\boldsymbol{p}^{(t+1)}) \max_{\boldsymbol{q} \in \Delta_m} h_{ij}(\boldsymbol{q}, 1) \\
& \leq 2 \frac{\max_{\boldsymbol{q} \in \Delta_m} h_{ij}(\boldsymbol{q}, 1)}{\min_{\substack{\boldsymbol{p} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} h_{ik}(\boldsymbol{p}, 1)} \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)}{h_{ik}(\boldsymbol{p}, 1)} \right\} \\
& = 2 \left(\max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)}{h_{ik}(\boldsymbol{p}, 1)} \right\} \right) \left(\max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)}{h_{ik}(\boldsymbol{p}, 1)} \right\} \right) \\
& = 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)}{h_{ik}(\boldsymbol{p}, 1)} \right\}^2 \\
& \leq v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) + 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)^2}{h_{ik}(\boldsymbol{p}, 1)^2} \right\}
\end{aligned}$$

Hence, the inductive hypothesis holds for $t + 1$. Putting it all together, we have, for all $t \in \mathbb{N}$:

$$v_i(\boldsymbol{p}^{(t)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) \leq v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) + 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)^2}{h_{ik}(\boldsymbol{p}, 1)^2} \right\}$$

Combining Lemma 6.5.1, and Lemma 6.5.2 with Theorem 6.1.1, we obtain our main result, namely a worst-case convergence rate of $O(^{(1+\epsilon^2)}/_t)$ for entropic *tâtonnement* in homothetic Fisher markets.

Theorem 6.5.1 [Convergence of Entropic Tâtonnement in Homothetic Fisher Markets]**.** Suppose $(\boldsymbol{u}, \boldsymbol{b})$ is a homothetic Fisher market and $\epsilon = \max_{\boldsymbol{p} \in \Delta_m, j, k \in [m]} \left| \epsilon_{h_{ij}, p_k}(\boldsymbol{p}, 1) \right|$. Then, the following holds for entropic *tâtonnement*: for all $t \in \mathbb{N}$,

$$\psi(\boldsymbol{p}^{(t)}) - \psi(\boldsymbol{p}^*) \leq \frac{\gamma \mathrm{div}_{\mathrm{KL}}(\boldsymbol{p}^*, \boldsymbol{p}^{(0)})}{t} \quad , \tag{6.13}$$

where $\gamma = \left(1 + \max_{j \in [m]} \sum_{i \in [n]} \left[v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{q} \in \Delta_m} h_{ij}(\boldsymbol{q}, 1) + 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m] : h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q},1)^2}{h_{ik}(\boldsymbol{p},1)^2} \right\} \right] \right) \left(6 + \frac{85\epsilon}{12} + \frac{25\epsilon^2}{72}\right)$.

Chapter 7

Appendix for Part I

7.1 Details of Section 5.4.3 Experiments

7.1.1 Computational Resources

Our experiments were run on MacOS machine with 8GB RAM and an Apple M1 chip, and took about 10 minutes to run. Only CPU resources were used.

7.1.2 Programming Languages, Packages, and Licensing

We ran our experiments in Python 3.7 (Van Rossum and Drake Jr, 1995), using NumPy (Harris et al., 2020), Jax (Bradbury et al., 2018), and JaxOPT (Blondel et al., 2021). All figures were graphed using Matplotlib (Hunter, 2007).

Python software and documentation are licensed under the PSF License Agreement. Numpy is distributed under a liberal BSD license. Pandas is distributed under a new BSD license. Matplotlib only uses BSD compatible code, and its license is based on the PSF license.

7.1.3 Experimental Setup Details

Each economy is initialized using a random seed to ensure reproducibility. Each consumer is assigned an initial endowment, drawn from a uniform distribution: $e' \sim \text{Unif}(10^{-6}, 1), \quad \forall i \in [n], j \in [m]$. For numerical stability, we restrict the total economy-wide aggregate supply of each commodity to remain fixed at 10[1], to this end we normalize the endowments of consumers for all $j \in [m]$, $i \in [n]$ to obtain their final endowment:

$$e_{ij} \doteq \frac{10e'_{ij}}{\sum_{i \in [n]} e'_{ij}}.$$

Each consumer's valuation of each commodity is drawn from a uniform distribution, i.e., for all $j \in [m]$, $i \in [n]$:

$$v_{ij} \sim \mathrm{Unif}(0, 1).$$

For any CES consumer $i \in [n]$, the elasticity of substitution parameter ρ_i, is drawn as follows from the uniform distribution for substitutes and complements consumers respectively:

$$\rho_i^{\text{substitutes}} \sim \mathrm{Unif}(0.6, 0.9) \qquad\qquad \rho_i^{\text{complements}} \sim \mathrm{Unif}(-1000, -1)$$

The initial price vector $\boldsymbol{p}^{(0)}$ for the algorithms is drawn from a uniform distribution s.t. for all $j \in [m]$:

$$p_j^{(0)} \sim \mathrm{Unif}(1, 10).$$

We note that while we initialize the prices between 1 and 10 for numerical stability, this choice is without loss of generality since the excess demand is homogeneous of degree 0.

To summarize. Given a random seed, the initialization process consists of: 1) Sampling endowments from a uniform distribution and normalizing them to ensure total supply constraints; 2) sampling valuations from a uniform distribution; 3) sampling substitution parameters for CES consumers, 4) generating an initial price vector.

7.2 Omitted Results and Proofs from Chapter 6

Lemma 7.2.1.

Suppose that u_i is homogeneous, i.e., $\forall \lambda > 0, u_i(\lambda \boldsymbol{x}_i) = \lambda u_i(\boldsymbol{x}_i)$. Then, the expenditure function and the Hicksian demand are homogeneous in ν_i, i.e., for all $\forall \lambda > 0$, $e_i(\boldsymbol{p}, \lambda\nu_i) = \lambda e_i(\boldsymbol{p}, \nu_i)$ and $\boldsymbol{h}_i(\boldsymbol{p}, \lambda\nu_i) = \lambda \boldsymbol{h}_i(\boldsymbol{p}, \nu_i)$. Likewise, the indirect utility function and

[1] This is without loss of generality since commodities are divisible.

the Marshallian demand are homogeneous in b_i, i.e., for all $\forall\lambda > 0$, $v_i(\boldsymbol{p}, \lambda b_i) = \lambda v_i(\boldsymbol{p}, b_i)$ and $\boldsymbol{d}_i(\boldsymbol{p}, \lambda b_i) = \lambda \boldsymbol{d}_i(\boldsymbol{p}, b_i)$.

Lemma 7.2.1

Without loss of generality, assume u_i is homogeneous of degree 1.[a]

For Hicksian demand, we have that:

$$
\begin{aligned}
&\boldsymbol{h}_i(\boldsymbol{p}, \lambda\nu_i) && (7.1)\\
&= \operatorname*{arg\,min}_{\boldsymbol{x}_i : u_i(\boldsymbol{x}_i) \geq \lambda\nu_i} \boldsymbol{p} \cdot \left(\lambda \frac{\boldsymbol{x}_i}{\lambda}\right) && (7.2)\\
&= \lambda \operatorname*{arg\,min}_{\boldsymbol{x}_i : u_i(\frac{\boldsymbol{x}_i}{\lambda}) \geq \nu_i} \boldsymbol{p} \cdot \left(\frac{\boldsymbol{x}_i}{\lambda}\right) && (7.3)\\
&= \operatorname*{arg\,min}_{\boldsymbol{x}_i : u_i(\boldsymbol{x}_i) \geq \nu_i} \boldsymbol{p} \cdot \boldsymbol{x}_i && (7.4)\\
&= \lambda \boldsymbol{h}_i(\boldsymbol{p}, \nu_i) \quad . && (7.5)
\end{aligned}
$$

The first equality follows from the definition of Hicksian demand; the second, by the homogeneity of u_i; the third, by the nature of constrained optimization; and the last, from the definition of Hicksian demand again. This result implies homogeneity of the expenditure function in ν_i:

$$e_i(\boldsymbol{p}, \lambda\nu_i) = \boldsymbol{h}_i(\boldsymbol{p}, \lambda\nu_i) \cdot \boldsymbol{p} = \lambda \boldsymbol{h}_i(\boldsymbol{p}, \nu_i) \cdot \boldsymbol{p} = \lambda e_i(\boldsymbol{p}, \nu_i) \quad .$$

The first and last equalities follow from the definition of the expenditure function, while the second equality follows from the homogeneity of Hicksian demand (Equation (7.5)).

The proof in the case of Marshallian demand and the indirect utility function is analogous.

[a] If the utility function is homogeneous of degree k, we can use a monotonic transformation, namely take the k^{th} root, to transform the utility function into one of degree 1, while still preserving the preferences that it represents.

Lemma 6.2.1.

If u_i is continuous and homogeneous of degree 1, then $v_i(\boldsymbol{p}, b_i)$ and $e_i(\boldsymbol{p}, \nu_i)$ are differentiable in b_i and ν_i, resp. Further, $\mathcal{D}_{b_i} v_i(\boldsymbol{p}, b_i) = \{v_i(\boldsymbol{p}, 1)\}$ and $\mathcal{D}_{\nu_i} e_i(\boldsymbol{p}, \nu_i) = \{e_i(\boldsymbol{p}, 1)\}$.

Lemma 6.2.1

We prove differentiability from first principles:

$$\begin{aligned}
\lim_{h \to 0} \frac{e_i(\boldsymbol{p}, \nu_i + h) - e_i(\boldsymbol{p}, \nu_i)}{h} &= \lim_{h \to 0} \frac{e_i(\boldsymbol{p}, (1)(\nu_i + h)) - e_i(\boldsymbol{p}, (1)\nu_i)}{h} \\
&= \lim_{h \to 0} \frac{e_i(\boldsymbol{p}, 1)(\nu_i + h) - e_i(\boldsymbol{p}, 1)(\nu_i)}{h} \\
&= \lim_{h \to 0} \frac{e_i(\boldsymbol{p}, 1)(\nu_i + h - \nu_i)}{h} \\
&= \lim_{h \to 0} \frac{e_i(\boldsymbol{p}, 1)(h)}{h} \\
&= e_i(\boldsymbol{p}, 1)
\end{aligned}$$

The first line follows from the definition of the derivative; the second line, by homogeneity of the expenditure function (Lemma 7.2.1), since u_i is homogeneous; and the final line follows from the properties of limits. The other two lines follow by simple algebra.

Hence, as $e_i(\boldsymbol{p}, \nu_i)$ is differentiable in ν_i, its subdifferential is a singleton with $\mathcal{D}_{\nu_i} e_i(\boldsymbol{p}, \nu_i) = \{e_i(\boldsymbol{p}, 1)\}$. The proof of the analogous result for the indirect utility function's derivative with respect to b_i is similar .

Corollary 6.2.1.

If buyer i's utility function u_i is CCH, then

$$\frac{1}{e_i(\boldsymbol{p}, 1)} = \frac{1}{\frac{\partial e_i(\boldsymbol{p}, \nu_i)}{\partial \nu_i}} = \frac{\partial v_i(\boldsymbol{p}, b_i)}{\partial b_i} = v_i(\boldsymbol{p}, 1) \quad . \tag{6.6}$$

Proof of Corollary 6.2.1

By Lemma 6.2.1, we know that $e_i(\boldsymbol{p}, \nu_i)$ is differentiable in ν_i and that $\mathcal{D}_{\nu_i} e_i(\boldsymbol{p}, \nu_i) = \{e_i(\boldsymbol{p}, 1)\}$. Similarly, by Lemma 6.2.1, we know that $\mathcal{D}_{b_i} v_i(\boldsymbol{p}, b_i)$ is differentiable in b_i

and that $\mathcal{D}_{b_i} v_i(\boldsymbol{p}, b_i) = \{v_i(\boldsymbol{p}, 1)\}$. Combining these facts yields:

$$\begin{aligned}
\mathcal{D}_{\nu_i} e_i(\boldsymbol{p}, \nu_i) \cdot \mathcal{D}_{b_i} v_i(\boldsymbol{p}, b_i) &= e_i(\boldsymbol{p}, 1) \cdot v_i(\boldsymbol{p}, 1) && \text{(Lemma 6.2.1)} \\
&= e_i(\boldsymbol{p}, v_i(\boldsymbol{p}, 1)) && \text{(Lemma 7.2.1)} \\
&= 1 && \text{(Equation (7.12))}
\end{aligned}$$

Therefore, $\frac{1}{\frac{\partial e_i(\boldsymbol{p}, \nu_i)}{\partial \nu_i}} = \frac{\partial v_i(\boldsymbol{p}, b_i)}{\partial b_i}$. Combining this conclusion with Lemma 6.2.1, we obtain the result.

Lemma 7.2.2.

Given a CCH Fisher market $(\boldsymbol{u}, \boldsymbol{b})$, the dual of our convex program (Theorem 6.2.1) and that of Eisenberg Gale differ by a constant, namely $\sum_{i \in [n]} (b_i \log b_i - b_i)$. In particular,

$$\begin{aligned}
&\min_{\boldsymbol{p} \in \mathbb{R}^m_+} \left\{ \sum_{j \in [m]} p_j - \sum_{i \in [n]} b_i \log\left(\partial_{\nu_i} e_i(\boldsymbol{p}, \nu_i)\right) \right\} \\
&= \min_{\boldsymbol{p} \in \mathbb{R}^m_+} \sum_{j \in [m]} p_j + \sum_{i \in [n]} \left(b_i \log\left(v_i(\boldsymbol{p}, b_i)\right) - b_i\right) - \sum_{i \in [n]} \left(b_i \log b_i - b_i\right)
\end{aligned}$$

Lemma 7.2.2

$$\begin{aligned}
&\min_{\boldsymbol{p} \in \mathbb{R}^m_+} \sum_{j \in [m]} p_j + \sum_{i \in [n]} \left(b_i \log\left(v_i(\boldsymbol{p}, b_i)\right) - b_i\right) \\
&= \min_{\boldsymbol{p} \in \mathbb{R}^m_+} \sum_{j \in [m]} p_j + \sum_{i \in [n]} b_i \log\left(b_i v_i(\boldsymbol{p}, 1)\right) - \sum_{i \in [n]} b_i && \text{(Lemma 7.2.1)} \\
&= \min_{\boldsymbol{p} \in \mathbb{R}^m_+} \left\{ \sum_{j \in [m]} p_j + \sum_{i \in [n]} b_i \log\left(v_i(\boldsymbol{p}, 1)\right) \right\} + \sum_{i \in [n]} b_i \log b_i - \sum_{i \in [n]} b_i \\
&= \min_{\boldsymbol{p} \in \mathbb{R}^m_+} \left\{ \sum_{j \in [m]} p_j - \sum_{i \in [n]} b_i \log\left(\frac{1}{v_i(\boldsymbol{p}, 1)}\right) \right\} + \sum_{i \in [n]} b_i \log b_i - \sum_{i \in [n]} b_i \\
&= \min_{\boldsymbol{p} \in \mathbb{R}^m_+} \left\{ \sum_{j \in [m]} p_j - \sum_{i \in [n]} b_i \log\left(e_i(\boldsymbol{p}, 1)\right) \right\} + \sum_{i \in [n]} b_i \log b_i - \sum_{i \in [n]} b_i && \text{(Corollary 6.2.1)} \\
&= \min_{\boldsymbol{p} \in \mathbb{R}^m_+} \left\{ \sum_{j \in [m]} p_j - \sum_{i \in [n]} b_i \log\left(\partial_{\nu_i} e_i(\boldsymbol{p}, \nu_i)\right) \right\} + \sum_{i \in [n]} b_i \log b_i - \sum_{i \in [n]} b_i && \text{(Lemma 6.2.1)}
\end{aligned}$$

Theorem 6.2.1 [New Convex Program for Homothetic Fisher Markets].

The optimal solutions $(\boldsymbol{X}^*, \boldsymbol{p}^*)$ to the following primal and dual convex programs correspond to Fisher equilibrium allocations and prices, respectively, of the homothetic Fisher market $(\boldsymbol{u}, \boldsymbol{b})$:

Primal

$$\max_{\boldsymbol{X} \in \mathbb{R}_+^{n \times m}} \quad \sum_{i \in [n]} \left[b_i \log u_i \left(\frac{\boldsymbol{x}_i}{b_i} \right) + b_i \right]$$

$$\text{subject to} \quad \sum_{i \in [n]} x_{ij} \leq 1 \qquad \forall j \in [m]$$

Dual

$$\min_{\boldsymbol{p} \in \Delta_m} \psi(\boldsymbol{p}) \doteq \sum_{j \in [m]} p_j - \sum_{i \in [n]} b_i \log \left(e_i(\boldsymbol{p}, 1) \right)$$

Theorem 6.2.1

By Lemma 7.2.2, our dual and the Eisenberg-Gale dual differ by a constant, which is independent of the decision variables $\boldsymbol{p} \in \mathbb{R}_+^m$. Hence, the optimal prices $\boldsymbol{p}^*$ of our dual are the same as those of the Eisenberg-Gale dual, and thus correspond to equilibrium prices in the CCH Fisher market $(\boldsymbol{u}, \boldsymbol{b})$. Finally, the objective function of our convex program's primal is:

$$\sum_{i \in [n]} b_i \log \left(u_i \left(\boldsymbol{x}_i \right) \right) - \sum_{i \in [n]} \left(b_i \log b_i - b_i \right) = \sum_{i \in [n]} b_i \log u_i \left(\frac{\boldsymbol{x}_i}{b_i} \right) + \sum_{i \in [n]} b_i \ .$$

Danskin's theorem (Danskin, 1966) offers insights into optimization problems of the form: $\min_{\boldsymbol{x} \in X} f(\boldsymbol{x}, \boldsymbol{p})$, where $X \subset \mathbb{R}^m$ is compact and non-empty. Among other things, Danskin's theorem allows us to compute the subdifferential of value of this optimization problem with respect to $\boldsymbol{p}$.

Theorem 7.2.1 [Danskin's Theorem (Danskin, 1966)].

Consider an optimization problem of the form: $\min_{\boldsymbol{x} \in X} f(\boldsymbol{x}, \boldsymbol{p})$, where $X \subset \mathbb{R}^m$ is compact and non-empty. Suppose that X is convex and that f is concave in $\boldsymbol{x}$. Let $V(\boldsymbol{p}) = \min_{\boldsymbol{x} \in X} f(\boldsymbol{x}, \boldsymbol{p})$ and $X^*(\boldsymbol{p}) = \arg\min_{\boldsymbol{x} \in X} f(\boldsymbol{x}, \boldsymbol{p})$. Then the subdifferential of V at $\widehat{\boldsymbol{p}}$ is given by $\mathcal{D}_{\boldsymbol{p}} V(\widehat{\boldsymbol{p}}) = \left\{ \nabla_{\boldsymbol{p}} f(\boldsymbol{x}^*(\widehat{\boldsymbol{p}}), \widehat{\boldsymbol{p}}) \mid \boldsymbol{x}^*(\widehat{\boldsymbol{p}}) \in X^*(\widehat{\boldsymbol{p}}) \right\}$.

Lemma 6.3.1 [Shephard's lemma, generalized for set-valued Hicksian demand (Blume, 2017; Shephard, 2015; Tanaka, 2008)].

Let $e_i(\boldsymbol{p}, \nu_i)$ be the expenditure function of buyer i and $\boldsymbol{h}_i(\boldsymbol{p}, \nu_i)$ be the Hicksian demand set of buyer i. The subdifferential $\mathcal{D}_{\boldsymbol{p}} e_i(\boldsymbol{p}, \nu_i)$ is the Hicksian demand at prices $\boldsymbol{p}$ and utility level ν_i, i.e., $\mathcal{D}_{\boldsymbol{p}} e_i(\boldsymbol{p}, \nu_i) = \boldsymbol{h}_i(\boldsymbol{p}, \nu_i)$.

Lemma 6.3.1

Recall that $e_i(\boldsymbol{p}, \nu_i) = \min_{\boldsymbol{x} \in \mathbb{R}_+^m : u_i(\boldsymbol{x}) \geq \nu_i} \boldsymbol{p} \cdot \boldsymbol{x}$. Without loss of generality, we can assume that consumption set is bounded from above, since utilities are assumed to represent locally non-satiated preferences, i.e., $\min_{\boldsymbol{x} \in X : u_i(\boldsymbol{x}) \geq \nu_i} \boldsymbol{p} \cdot \boldsymbol{x}$ where $X \subset \mathbb{R}_+^m$ is compact. Using Danskin's theorem:

$$\begin{aligned}
\mathcal{D}_{\boldsymbol{p}} e_i(\boldsymbol{p}, \nu_i) &= \left\{ \nabla_{\boldsymbol{p}} \left(\boldsymbol{p} \cdot \boldsymbol{x} \right) \left(\boldsymbol{x}^*(\boldsymbol{p}, \nu_i) \right) \mid \boldsymbol{x}^*(\boldsymbol{p}, \nu_i) \in \boldsymbol{h}_i(\boldsymbol{p}, \nu_i) \right\} && \text{(Danskin's Thm)} \\
&= \left\{ \boldsymbol{x}^*(\boldsymbol{p}, \nu_i) \mid \boldsymbol{x}^*(\boldsymbol{p}, \nu_i) \in \boldsymbol{h}_i(\boldsymbol{p}, \nu_i) \right\} \\
&= \boldsymbol{h}_i(\boldsymbol{p}, \nu_i)
\end{aligned}$$

The first equality follows from Danskin's theorem, using the facts that the objective of the expenditure minimization problem is affine and the constraint set is compact. The second equality follows by calculus, and the third, by the definition of Hicksian demand.

Theorem 6.3.1.

Given any homothetic Fisher market $(\boldsymbol{u}, \boldsymbol{b})$, the subdifferential of the dual of the program in Theorem 6.2.1 at any price $\boldsymbol{p}$ is equal to the negative excess demand in $(\boldsymbol{u}, \boldsymbol{b})$ at price $\boldsymbol{p}$: i.e., $\mathcal{D}_{\boldsymbol{p}} \psi(\boldsymbol{p}) = -\mathcal{Z}(\boldsymbol{p})$.

Theorem 6.3.1

For all goods $j \in [m]$, we have:

$$
\begin{aligned}
&\mathcal{D}_{p_j}\left(\sum_{j\in[m]} p_j - \sum_{i\in[n]} b_i \log \partial_{\nu_i} e_i(\boldsymbol{p}, \nu_i)\right) \\
&= \{1\} - \mathcal{D}_{p_j}\left(\sum_{i\in[n]} b_i \log \partial_{\nu_i} e_i(\boldsymbol{p}, \nu_i)\right) \\
&= \{1\} - \sum_{i\in[n]} \mathcal{D}_{p_j}\left(b_i \log \partial_{\nu_i} e_i(\boldsymbol{p}, \nu_i)\right) \\
&= \{1\} - \sum_{i\in[n]} d_{ij}(\boldsymbol{p}, b_i) && \text{(Lemma 6.3.2)} \\
&= -z_j(\boldsymbol{p})
\end{aligned}
$$

Lemma 6.3.2.

If buyer i's utility function u_i is continuous and homogeneous, then $\mathcal{D}_{\boldsymbol{p}}\left(b_i \log\left(\frac{\partial e_i(\boldsymbol{p},\nu_i)}{\partial \nu_i}\right)\right) = \boldsymbol{d}_i(\boldsymbol{p}, b_i)$.

Lemma 6.3.2

Without loss of generality, we can assume u_i is homogeneous of degree 1. Then:

$$
\begin{aligned}
\mathcal{D}_{\boldsymbol{p}}\left(b_i \log\left(\frac{\partial e_i(\boldsymbol{p},\nu_i)}{\partial \nu_i}\right)\right) &= \left(\frac{b_i}{\frac{\partial e_i(\boldsymbol{p},\nu_i)}{\partial \nu_i}}\right) \mathcal{D}_{\boldsymbol{p}}\left(\frac{\partial e_i(\boldsymbol{p},\nu_i)}{\partial \nu_i}\right) \\
&= b_i\left(\frac{\partial v_i(\boldsymbol{p},b_i)}{\partial b_i}\right) \mathcal{D}_{\boldsymbol{p}}\left(\frac{\partial e_i(\boldsymbol{p},\nu_i)}{\partial \nu_i}\right) && \text{(Corollary 6.2.1)} \\
&= b_i\left(\frac{\partial v_i(\boldsymbol{p},b_i)}{\partial b_i}\right) \mathcal{D}_{\boldsymbol{p}} e_i(\boldsymbol{p}, 1) && \text{(Lemma 6.2.1)} \\
&= b_i\left(\frac{\partial v_i(\boldsymbol{p},b_i)}{\partial b_i}\right) \boldsymbol{h}_i(\boldsymbol{p}, 1) && \text{(Shephard's Lemma)} \\
&= b_i\, v_i(\boldsymbol{p}, 1)\, \boldsymbol{h}_i(\boldsymbol{p}, 1) && \text{(Lemma 6.2.1)} \\
&= v_i(\boldsymbol{p}, b_i)\, \boldsymbol{h}_i(\boldsymbol{p}, 1) && \text{(Lemma 7.2.1)} \\
&= \boldsymbol{h}_i\left(\boldsymbol{p}, v_i(\boldsymbol{p}, b_i)\right) && \text{(Lemma 7.2.1)} \\
&= \boldsymbol{d}_i(\boldsymbol{p}, b_i) && \text{(Equation (7.14))}
\end{aligned}
$$

We start by presenting the first lemma, which shows that the utility level elasticity of Hicksian demand is equal to 1 in homothetic Fisher markets.

Lemma 6.4.1.

For any Hicksian demand $\boldsymbol{h}_i$ associated with a homogeneous utility function u_i, for all $j, k \in [m], \boldsymbol{p} \in \mathbb{R}_+^m, \nu_i \in \mathbb{R}_+$, it holds that $\epsilon_{h_{ij},p_k}(\boldsymbol{p}, \nu_i) = \epsilon_{h_{ij},p_k}(\boldsymbol{p}, 1) = 1$.

Proof of Lemma 6.4.1

Recall from Goktas et al. (2022b) that for homogeneous utility functions, the Hicksian demand is homogeneous in ν, i.e., for all $\lambda \geq 0$, $\boldsymbol{h}_i(\boldsymbol{p}, \lambda\nu) = \lambda \boldsymbol{h}_i(\boldsymbol{p}, \nu)$. Hence, we have:

$$\epsilon_{h_{ij},\nu_i}(\boldsymbol{p}, \nu_i) = \mathcal{D}_{\nu_i} h_{ij}(\boldsymbol{p}, \nu_i) \frac{\nu_i}{h_{ij}(\boldsymbol{p}, \nu_i)} \tag{7.6}$$

$$= \nu_i \mathcal{D}_{\nu_i} h_{ij}(\boldsymbol{p}, 1) \frac{\nu_i}{\nu_i h_{ij}(\boldsymbol{p}, 1)} \quad \text{(Homogeneity of Hicksian demand)} \tag{7.7}$$

$$= \frac{\nu_i}{h_{ij}(\boldsymbol{p}, 1)} \mathcal{D}_{\nu_i} h_{ij}(\boldsymbol{p}, 1) \tag{7.8}$$

$$= \frac{\nu_i}{h_{ij}(\boldsymbol{p}, 1)} \mathcal{D}_{\nu_i} h_{ij}(\boldsymbol{p}, 1) \tag{7.9}$$

$$= \epsilon_{h_{ij},\nu_i}(\boldsymbol{p}, 1) \tag{7.10}$$

Additionally, looking back at Equation (7.9), since Hicksian demand is homogeneous of degree 1 in ν_i for homogeneous utility function (Lemma 7.2.1), by Euler's theorem for homogeneous functions (see, for instance, (Border, 2017)), we have: $\frac{\nu_i}{h_{ij}(\boldsymbol{p},1)} \mathcal{D}_{\nu_i} h_{ij}(\boldsymbol{p}, 1) = \frac{h_{ij}(\boldsymbol{p},1)}{h_{ij}(\boldsymbol{p},1)} = 1$.

We recall Shephard's lemma which was used in the Equation (6.8):

Lemma 6.3.1 [Shephard's lemma, generalized for set-valued Hicksian demand (Blume, 2017; Shephard, 2015; Tanaka, 2008)]**.**

Let $e_i(\boldsymbol{p}, \nu_i)$ be the expenditure function of buyer i and $\boldsymbol{h}_i(\boldsymbol{p}, \nu_i)$ be the Hicksian demand set of buyer i. The subdifferential $\mathcal{D}_{\boldsymbol{p}} e_i(\boldsymbol{p}, \nu_i)$ is the Hicksian demand at prices $\boldsymbol{p}$ and utility level ν_i, i.e., $\mathcal{D}_{\boldsymbol{p}} e_i(\boldsymbol{p}, \nu_i) = \boldsymbol{h}_i(\boldsymbol{p}, \nu_i)$.

We first prove that by setting γ to be 5 times the maximum demand for any good throughout the entropic *tâtonnement* process , we can bound the change in the prices of goods in each round. We will use the fact that the change in the price of each good is bounded as an assumption in most of the following results.

Lemma 7.2.3.

Suppose that entropic *tâtonnement* process is run for all $t \in [T] \subseteq \mathbb{N}_+$ with $\gamma \geq 5 \max_{\substack{t \in [T] \\ j \in [m]}} \{1, q_j^{(t)}\}$ and let $\Delta \boldsymbol{p} = \boldsymbol{p}^{(t+1)} - \boldsymbol{p}^{(t)}$. then the following holds for all $t \in \mathbb{N}$:

$$e^{-\frac{1}{5}} p_j^{(t)} \leq p_j^{(t+1)} \leq e^{\frac{1}{5}} p_j^{(t)} \text{ and } \frac{|\Delta p_j|}{p_j^{(t)}} \leq \frac{1}{4}$$

Lemma 7.2.3

The price of of a good $j \in [m]$ can at most increase by a factor of $e^{\frac{1}{5}}$:

$$p_j^{(t+1)} = p_j^{(t)} e^{\frac{z_j(\boldsymbol{p}^{(t)})}{\gamma}} = p_j^{(t)} \exp\left\{\frac{q_j^{(t)} - 1}{\gamma}\right\} \leq p_j^{(t)} \exp\left\{\frac{q_j^{(t)}}{\gamma}\right\} \leq p_j^{(t)} \exp\left\{\frac{q_j^{(t)}}{5 \max_{\substack{t \in \mathbb{N} \\ j \in [m]}} \{1, q_j^{(t)}\}}\right\} \leq p_j^{(t)} e^{\frac{1}{5}}$$

and decrease by a factor of $e^{-\frac{1}{5}}$:

$$p_j^{(t+1)} = p_j^{(t)} e^{\frac{z_j(\boldsymbol{p}^{(t)})}{\gamma}} = p_j^{(t)} \exp\left\{\frac{q_j^{(t)} - 1}{\gamma}\right\} \geq p_j^{(t)} \exp\left\{\frac{-1}{\gamma}\right\} \geq p_j^{(t)} \exp\left\{\frac{-1}{5 \max_{\substack{t \in \mathbb{N} \\ j \in [m]}} \{1, q_j^{(t)}\}}\right\} \geq p_j^{(t)} e^{-\frac{1}{5}}$$

Hence, we have $e^{-\frac{1}{5}} p_j^{(t)} \leq p_j^{(t+1)} \leq e^{\frac{1}{5}} p_j^{(t)}$. Substracting $p_j^{(t)}$ from both sides and dividing by $p_j^{(t)}$, we obtain:

$$\frac{|\Delta p_j|}{p_j^{(t)}} = \frac{|p_j^{(t+1)} - p_j^{(t)}|}{p_j^{(t)}} \leq e^{1/5} - 1 \leq \frac{1}{4}$$

The following two results are due to Cheung et al. (2013). We include their proofs for completeness. They allows us to relate the change in prices to the KL-divergence.

Lemma 7.2.4 [Cheung et al. (2013)]**.**

Fix $t \in \mathbb{N}_+$ and let $\Delta \boldsymbol{p} = \boldsymbol{p}^{(t+1)} - \boldsymbol{p}^{(t)}$. Suppose that for all $j \in [m]$, $\frac{|\Delta p_j|}{p_j^{(t)}} \leq \frac{1}{4}$, then:

$$\frac{(\Delta p_j)^2}{p_j^{(t)}} \leq \frac{9}{2} \mathrm{div}_{\mathrm{KL}}(p_j^{(t)} + \Delta p_j, p_j^{(t)}) \tag{7.11}$$

Lemma 7.2.4

The bound $\log(x) \geq x - x^2$ for $|x|\leq \frac{1}{4}$ is used below:

$$\begin{aligned}
&\mathrm{div}_{\mathrm{KL}}(p_j^{(t)} + \Delta p_j, p_j^{(t)}) \\
&= (p_j^{(t)} + \Delta p_j)(\log(p_j^{(t)} + \Delta p_j)) - (p_j^{(t)} + \Delta p_j - p_j^{(t)} \log(p_j) + p_j^{(t)} - \log(p_j^{(t)})\Delta p_j \\
&= -\Delta p_j + (p_j^{(t)} + \Delta p_j) \log \left(1 + \frac{\Delta p_j}{p_j^{(t)}}\right) \\
&\geq -\Delta p_j + (p_j^{(t)} + \Delta p_j) \left(\frac{\Delta p_j}{p_j^{(t)}} - \frac{11}{18} \frac{(\Delta p_j)^2}{(p_j^{(t)})^2}\right) \\
&\geq \frac{7}{18} \frac{(\Delta p_j)^2}{p_j^{(t)}} \left(1 - \frac{11}{7} \frac{\Delta p_j}{p_j^{(t)}}\right) \\
&= \frac{7}{18} \frac{17}{28} \frac{(\Delta p_j)^2}{p_j^{(t)}} \\
&\geq \frac{2}{9} \frac{(\Delta p_j)^2}{p_j^{(t)}}
\end{aligned}$$

Lemma 7.2.5.

Fix $t \in \mathbb{N}_+$ and let $\Delta \boldsymbol{p} = \boldsymbol{p}^{(t+1)} - \boldsymbol{p}^{(t)}$. Suppose that $\frac{|\Delta p_j|}{p_j} \leq \frac{1}{4}$, then for any $c \in (0, 1)$, and $\boldsymbol{A} \in \mathbb{R}^{n \times m}$, and for all $j \in [m]$:

$$\frac{1}{b_i} \sum_{j \in [m]} \sum_{k \in [m]} a_{il} d_{ik}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)|\Delta p_j||\Delta p_k|\leq \frac{4}{3} \sum_{l \in [m]} \frac{a_{il}}{p_l^{(t)}} (\Delta p_l)^2$$

Lemma 7.2.5

First, note that since by our assumption the utilities are locally non-satiated, Walras' law is satisfied, i.e., we have $b_i = \sum_{k \in [m]} d_{ik}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)(p_k^{(t)} + c\Delta p_k)$;

$$b_i \sum_{l \in [m]} \frac{a_{il}}{p_l^{(t)}} (\Delta p_l)^2$$
$$= \sum_{l \in [m]} \frac{\left(\sum_{k \in [m]} d_{ik}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)(p_k^{(t)} + c\Delta p_k)\right) d_{il}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)}{p_l^{(t)}} (\Delta p_l)^2$$
$$\geq \sum_{l \in [m]} \frac{\left(\sum_{k \in [m]} d_{ik}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)(p_k^{(t)} - \frac{1}{4} p_k^{(t)})\right) a_{il}}{p_l^{(t)}} (\Delta p_l)^2$$
$$= \sum_{l \in [m]} \frac{\left(\sum_{k \in [m]} d_{ik}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)(\frac{3}{4} p_k^{(t)})\right) a_{il}}{p_l^{(t)}} (\Delta p_l)^2$$
$$= \frac{3}{4} \sum_{l \in [m]} \sum_{k \in [m]} a_{il} d_{il}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i) \frac{p_k^{(t)}}{p_l^{(t)}} (\Delta p_l)^2$$
$$= \frac{3}{4} \left[\sum_{l \in [m]} a_{il} d_{il}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)(\Delta p_l)^2 + \sum_{l \in [m]} \sum_{k \neq l} a_{il} d_{ik}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i) \frac{p_k^{(t)}}{p_l^{(t)}} (\Delta p_l)^2 \right]$$
$$= \frac{3}{4} \left[\sum_{l \in [m]} a_{il} d_{il}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)(\Delta p_l)^2 + \sum_{k \in [m]} \sum_{k \leq l} a_{il} d_{ik}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i) \left(\frac{p_k^{(t)}}{p_l^{(t)}} |\Delta p_l|^2 + \frac{p_l^{(t)}}{p_k^{(t)}} |\Delta p_k|^2 \right) \right]$$

Now, we apply the AM-GM inequality, i.e., for all $x, y \in \mathbb{R}_+$ since $\sqrt{xy} \leq \frac{x+y}{2}$, we have:

$$b_i \sum_{l \in [m]} \frac{a_{il}}{p_l^{(t)}} (\Delta p_l)^2$$
$$\geq \frac{3}{4} \sum_{l \in [m]} a_{il} d_{il}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)(\Delta p_l)^2 + \sum_{k < l} a_{il} d_{ik}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i) \left(2|\Delta p_l||\Delta p_k|\right)$$
$$= \frac{3}{4} \sum_{j \in [m]} \sum_{k \in [m]} a_{il} d_{ik}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)|\Delta p_j||\Delta p_k|$$

Lemma 7.2.6.

(Cheung et al., 2013) For all $j \in [m]$:

$$\frac{1}{b_i} \sum_{j \in [m]} \sum_{k \in [m]} d_{ij}^{(t)} d_{ik}^{(t)} |\Delta p_j||\Delta p_k| \leq \sum_{l \in [m]} \frac{d_{il}^{(t)}}{p_l^{(t)}} (\Delta p_l)^2$$

Lemma 7.2.6

First, note that by Walras' law we have $b_i = \sum_{k \in [m]} d_{ik}^{(t)} p_k^{(t)}$;

$$
\begin{aligned}
b_i \sum_{l \in [m]} \frac{d_{il}^{(t)}}{p_l^{(t)}} (\Delta p_l)^2 &= \sum_{l \in [m]} \frac{\left(\sum_{k \in [m]} d_{ik}^{(t)} p_k^{(t)}\right) d_{il}^{(t)}}{p_l^{(t)}} (\Delta p_l)^2 \\
&= \sum_{l \in [m]} \sum_{k \in [m]} d_{il}^{(t)} d_{ik}^{(t)} \frac{p_k^{(t)}}{p_l^{(t)}} (\Delta p_l)^2 \\
&= \sum_{l \in [m]} (d_{il}^{(t)})^2 (\Delta p_l)^2 + \sum_{l \in [m]} \sum_{k \neq l} d_{il}^{(t)} d_{ik}^{(t)} \frac{p_k^{(t)}}{p_l^{(t)}} (\Delta p_l)^2 \\
&= \sum_{l \in [m]} (d_{il}^{(t)})^2 (\Delta p_l)^2 + \sum_{k \in [m]} \sum_{k \leq l} d_{ik}^{(t)} d_{il}^{(t)} \left(\frac{p_k^{(t)}}{p_l^{(t)}} (\Delta p_l)^2 + \frac{p_l^{(t)}}{p_k^{(t)}} (\Delta p_k)^2 \right)
\end{aligned}
$$

Now, we apply the AM-GM inequality:

$$
\begin{aligned}
b_i \sum_{l \in [m]} \frac{d_{il}^{(t)}}{p_l^{(t)}} (\Delta p_l)^2 &\geq \sum_{l \in [m]} (d_{il}^{(t)})^2 (\Delta p_l)^2 + \sum_{k < l} d_{ik}^{(t)} d_{il}^{(t)} \left(2 |\Delta p_l| |\Delta p_k| \right) \\
&= \sum_{j \in [m]} \sum_{k \in [m]} d_{ij}^{(t)} d_{ik}^{(t)} |\Delta p_j| |\Delta p_k|
\end{aligned}
$$

An important result in microeconomics is the **law of demand** which states that when the price of a good increases, the Hicksian demand for that good decreases in a very general setting of utility functions (Levin, 2004; Mas-Colell et al., 1995). We state a weaker version of the law of demand which is re-formulated to fit the tâtonnement framework.

Lemma 7.2.7 [Law of Demand]**.**
(Levin, 2004; Mas-Colell et al., 1995) Suppose that $\forall j \in [m], t \in \mathbb{N}, p_j^{(t)}, p_j^{(t+1)} \geq 0$. Then, $\sum_{j \in [m]} \Delta p_j \left(h_{ij}^{(t+1)} - h_{ij}^{(t)} \right) \leq 0$.

A simple corollary of the law of demand which is used throughout the rest of this paper is that, during tâtonnement, the change in expenditure of the next time period is always less than or equal to the change in expenditure of the previous time period's.

Corollary 7.2.1.

Suppose that $\forall t \in \mathbb{N}, j \in [m], p_j^{(t)}, p_j^{(t+1)} \geq 0$, then $\forall t \in \mathbb{N}, \sum_{j \in [m]} \Delta p_j h_{ij}^{(t+1)} \leq \sum_{j \in [m]} \Delta p_j h_{ij}^{(t)}$.

The following lemma simply restates an essential fact about expenditure functions and Hicksian demand, namely that the Hicksian demand is the minimizer of the expenditure function.

Lemma 7.2.8.

For all $t \in \mathbb{N}$, we have $\sum_{j \in [m]} h_{ij}^{(t)} p_j^{(t)} \leq \sum_{j \in [m]} h_{ij}^{(t+1)} p_j^{(t)}$.

Lemma 7.2.8

For the sake of contradiction, assume that $\sum_{j \in [m]} h_{ij}^{(t)} p_j^{(t)} > \sum_{j \in [m]} h_{ij}^{(t+1)} p_j^{(t)}$. By the definition of the Hicksian demand, we know that the bundle $\boldsymbol{h}_i^{(t)}$ provides the buyer with one unit of utility. Recall that the expenditure at any price $\boldsymbol{p}$ is equal to the sum of the product of the Hicksian demands and prices, that is $e_i(\boldsymbol{p}, 1) = \sum_{j \in [m]} h_{ij}(\boldsymbol{p}, 1) p_j$. Hence, we have $e_i(p_j^{(t)}, 1) = \sum_{j \in [m]} h_{ij}^{(t)} p_j^{(t)} > \sum_{j \in [m]} h_{ij}^{(t+1)} p_j^{(t)} = e_i(\boldsymbol{p}^{(t)}, 1)$, a contradiction.

We now introduce the following lemma which makes use of results on the behavior of Hicksian demand and expenditure functions in homothetic Fisher markets introduced by Goktas et al. (2022b). In conjunction with Corollary 7.2.1 and Lemma 7.2.8 are key in proving that Lemma 6.5.1 holds allowing us to establish convergence of tâtonnement in a general setting of utility functions. Additionally, the lemma relates the Marshallian demand of homogeneous utility functions to their Hicksian demand. Before we present the lemma, we recall the following identities (Mas-Colell et al., 1995):

$$\forall b_i \in \mathbb{R}_+ \qquad e_i(\boldsymbol{p}, v_i(\boldsymbol{p}, b_i)) = b_i \tag{7.12}$$

$$\forall \nu_i \in \mathbb{R}_+ \qquad v_i(\boldsymbol{p}, e_i(\boldsymbol{p}, \nu_i)) = \nu_i \tag{7.13}$$

$$\forall b_i \in \mathbb{R}_+ \qquad \boldsymbol{h}_i(\boldsymbol{p}, v_i(\boldsymbol{p}, b_i)) = \boldsymbol{d}_i(\boldsymbol{p}, b_i) \tag{7.14}$$

$$\forall \nu_i \in \mathbb{R}_+ \qquad \boldsymbol{d}_i(\boldsymbol{p}, e_i(\boldsymbol{p}, \nu_i)) = \boldsymbol{h}_i(\boldsymbol{p}, \nu_i) \tag{7.15}$$

Lemma 7.2.9.

Suppose that u_i is continuous and homogeneous, then the following holds:

$$\forall j \in [m] \qquad d_{ij}(\boldsymbol{p}, b_i) = \frac{b_i h_{ij}(\boldsymbol{p}, 1)}{\sum_{j \in [m]} h_{ij}(\boldsymbol{p}, 1) p_j}$$

Lemma 7.2.9

We note that when utility function u_i is strictly concave, the Marshallian and Hicksian demand are unique making the following equalities well-defined.

$$\begin{aligned}
\frac{b_i h_{ij}(\boldsymbol{p}, 1)}{\sum_{j \in [m]} h_{ij}(\boldsymbol{p}, 1) p_j} &= \frac{b_i h_{ij}(\boldsymbol{p}, 1)}{e_i(\boldsymbol{p}, 1)} && \text{(Definition of expenditure function)} \\
&= b_i v_i(\boldsymbol{p}, 1) h_{ij}(\boldsymbol{p}, 1) && \text{(Corollary 1 of Goktas et al. (2022b))} \\
&= v_i(\boldsymbol{p}, b_i) h_{ij}(\boldsymbol{p}, 1) \\
&= h_{ij}(\boldsymbol{p}, v_i(\boldsymbol{p}, b_i)) \\
&= d_{ij}(\boldsymbol{p}, b_i) && \text{(Marshallian Identity Equation (7.14))}
\end{aligned}$$

The following lemma proves that the relative change in expenditures at each iteration of tatonnement is bounded when the relative change in prices is bounded.

Lemma 7.2.10.

Suppose that $\forall j \in [m], \frac{|\Delta p_j|}{p_j^{(t)}} \leq \frac{1}{4}$, then for any $t \in \mathbb{N}_+$ and $i \in [n]$:

$$\left| \frac{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t+1)}, 1), \boldsymbol{p}^{(t+1)} \right\rangle - \left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \boldsymbol{p}^{(t)} \right\rangle}{\left\langle \boldsymbol{h}_i^{(t)}, \boldsymbol{p}^{(t)} \right\rangle} \right| \leq \frac{1}{4} \tag{7.16}$$

Proof

Case 1: $\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t+1)}, 1), \boldsymbol{p}^{(t+1)} \right\rangle \geq \left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \boldsymbol{p}^{(t)} \right\rangle$

$$\frac{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t+1)}, 1), \boldsymbol{p}^{(t+1)} \right\rangle - \left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \boldsymbol{p}^{(t)} \right\rangle}{\left\langle \boldsymbol{h}_i^{(t)}, \boldsymbol{p}^{(t)} \right\rangle} \tag{7.17}$$

$$\leq \frac{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \boldsymbol{p}^{(t+1)} \right\rangle - \left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \boldsymbol{p}^{(t)} \right\rangle}{\left\langle \boldsymbol{h}_i^{(t)}, \boldsymbol{p}^{(t)} \right\rangle} \qquad \text{(Corollary 7.2.1)} \tag{7.18}$$

$$= \frac{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \boldsymbol{p}^{(t+1)} \right\rangle}{\left\langle \boldsymbol{h}_i^{(t)}, \boldsymbol{p}^{(t)} \right\rangle} - 1 \tag{7.19}$$

$$\leq \frac{5}{4} \frac{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \boldsymbol{p}^{(t)} \right\rangle}{\left\langle \boldsymbol{h}_i^{(t)}, \boldsymbol{p}^{(t)} \right\rangle} - 1 \tag{7.20}$$

$$= \frac{1}{4} \tag{7.21}$$

where the penultimate line follows from the assumption that $\forall j \in [m]$, $\frac{|\Delta p_j|}{p_j^{(t)}} \leq \frac{1}{4}$.

Case 2: $\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t+1)}, 1), \boldsymbol{p}^{(t+1)} \right\rangle \leq \left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \boldsymbol{p}^{(t)} \right\rangle$

$$\frac{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \boldsymbol{p}^{(t)} \right\rangle - \left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t+1)}, 1), \boldsymbol{p}^{(t+1)} \right\rangle}{\left\langle \boldsymbol{h}_i^{(t)}, \boldsymbol{p}^{(t)} \right\rangle} \tag{7.22}$$

$$= 1 - \frac{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t+1)}, 1), \boldsymbol{p}^{(t+1)} \right\rangle}{\left\langle \boldsymbol{h}_i^{(t)}, \boldsymbol{p}^{(t)} \right\rangle} \tag{7.23}$$

$$\leq 1 - \frac{3}{4} \frac{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t+1)}, 1), \boldsymbol{p}^{(t)} \right\rangle}{\left\langle \boldsymbol{h}_i^{(t)}, \boldsymbol{p}^{(t)} \right\rangle} \tag{7.24}$$

$$\leq 1 - \frac{3}{4} \frac{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \boldsymbol{p}^{(t)} \right\rangle}{\left\langle \boldsymbol{h}_i^{(t)}, \boldsymbol{p}^{(t)} \right\rangle} \qquad \text{(Corollary 7.2.1)} \tag{7.25}$$

$$= \frac{1}{4} \tag{7.26}$$

where the second line follows from the assumption that $\forall j \in [m]$, $\frac{|\Delta p_j|}{p_j^{(t)}} \leq \frac{1}{4}$.

Lemma 7.2.11.

Suppose that for all $j \in [m]$, $\frac{|\Delta p_j|}{p_j} \leq \frac{1}{4}$, then for some $c \in (0,1)$ and $t \in \mathbb{N}_+$, we have:

$$\frac{1}{b_i}\left(\left\langle \boldsymbol{d}_i^{(t)}, \Delta \boldsymbol{p}\right\rangle + \frac{b_i}{2}\frac{\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)}\right)^2 \tag{7.27}$$

$$\leq \left(1 + \frac{5\epsilon}{9}\right)\sum_{l\in[m]}\frac{d_{il}^{(t)}}{p_l^{(t)}}(\Delta p_l)^2 + \left(\frac{25\epsilon^2}{432}\right)\sum_{l\in[m]}\frac{d_{il}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, b_i)}{p_l^{(t)}}(\Delta p_l)^2 \tag{7.28}$$

Proof of Lemma 7.2.11

$$\begin{aligned}
&\frac{1}{b_i}\left(\left\langle \boldsymbol{d}_i^{(t)}, \Delta \boldsymbol{p}\right\rangle + \frac{b_i}{2}\frac{\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)}\right)^2 \\
&= \frac{1}{b_i}\Bigg[\left\langle \boldsymbol{d}_i^{(t)}, \Delta \boldsymbol{p}\right\rangle^2 + 2\left\langle \boldsymbol{d}_i^{(t)}, \Delta \boldsymbol{p}\right\rangle\left(\frac{b_i}{2}\frac{\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)}\right) \\
&\quad + \left(\frac{b_i}{2}\frac{\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)}\right)^2\Bigg] \\
&\leq \frac{1}{b_i}\Bigg[\left|\left\langle \boldsymbol{d}_i^{(t)}, \Delta \boldsymbol{p}\right\rangle^2\right| + 2\left|\left\langle \boldsymbol{d}_i^{(t)}, \Delta \boldsymbol{p}\right\rangle\right|\left|\frac{b_i}{2}\frac{\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)}\right| \\
&\quad + \left|\frac{b_i}{2}\frac{\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)}\right|^2\Bigg] \\
&\leq \frac{1}{b_i}\Bigg[\left\langle \boldsymbol{d}_i^{(t)}, |\Delta \boldsymbol{p}|\right\rangle^2 + 2\left\langle \boldsymbol{d}_i^{(t)}, |\Delta \boldsymbol{p}|\right\rangle\left|\frac{b_i}{2}\frac{\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)}\right| \\
&\quad + + \left|\frac{b_i}{2}\frac{\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)}\right|^2\Bigg]
\end{aligned}$$

where we denote $|\Delta\boldsymbol{p}| = (|\Delta p_1|, \ldots, |\Delta p_m|)$.

$$\begin{aligned}
&\leq \frac{1}{b_i}\Bigg[\left\langle \boldsymbol{d}_i^{(t)}, |\Delta \boldsymbol{p}|\right\rangle^2 + 2\left\langle \boldsymbol{d}_i^{(t)}, |\Delta \boldsymbol{p}|\right\rangle \\
&\left(\frac{5\epsilon}{6}\sum_j \frac{(\Delta p_j)^2}{p_j} d_{ij}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, b_i)\right) + \left(\frac{5\epsilon}{6}\sum_j \frac{(\Delta p_j)^2}{p_j} d_{ij}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, b_i)\right)^2\Bigg] \quad \text{(Lemma 6.4.2)} \\
&= \frac{1}{b_i}\Bigg[\left\langle \boldsymbol{d}_i^{(t)}, |\Delta \boldsymbol{p}|\right\rangle^2 + \frac{5\epsilon}{3}\left\langle \boldsymbol{d}_i^{(t)}, |\Delta \boldsymbol{p}|\right\rangle\left(\sum_j \frac{(\Delta p_j)^2}{p_j} d_{ij}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, b_i)\right) \\
&+ \frac{25\epsilon^2}{36}\left(\sum_j \frac{(\Delta p_j)^2}{p_j} d_{ij}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, b_i)\right)^2\Bigg]
\end{aligned}$$

Since $\forall j \in [m], \frac{|\Delta p_j|}{p_j^{(t)}} \leq \frac{1}{4}$, we have:

$$\leq \frac{1}{b_i}\left[\left\langle \boldsymbol{d}_i^{(t)}, |\Delta \boldsymbol{p}|\right\rangle^2 + \frac{5\epsilon}{12}\left\langle \boldsymbol{d}_i^{(t)}, |\Delta \boldsymbol{p}|\right\rangle \left(\sum_j |\Delta p_j|\, d_{ij}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, 1)\right)\right.$$

$$\left. + + \frac{25\epsilon^2}{576}\left(\sum_j |\Delta p_j| d_{ij}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)\right)^2\right] \tag{7.29}$$

$$= \frac{1}{b_i}\sum_{j\in[m]}\sum_{k\in[m]} d_{ij}^{(t)} d_{ik}^{(t)} |\Delta p_j||\Delta p_k| + \frac{1}{b_i}\frac{5\epsilon}{12}\sum_j\sum_k d_{ik}^{(t)} d_{ij}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)\,|\Delta p_k|\,|\Delta p_j|$$

$$+ \frac{1}{b_i}\frac{25\epsilon^2}{576}\sum_j\sum_k d_{ij}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i) d_{ik}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)|\Delta p_k||\Delta p_j| \tag{7.30}$$

$$\leq \sum_{l\in[m]} \frac{d_{il}^{(t)}}{p_l^{(t)}}(\Delta p_l)^2 + \frac{1}{b_i}\frac{5\epsilon}{12}\sum_j\sum_k d_{ik}^{(t)} d_{ij}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)\,|\Delta p_k|\,|\Delta p_j|$$

$$+ \frac{1}{b_i}\frac{25\epsilon^2}{576}\sum_j\sum_k d_{ij}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i) d_{ik}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)|\Delta p_k||\Delta p_j| \tag{7.31}$$

where the last line was obtained by (Lemma 7.2.6). Continuing, by Lemma 7.2.5, we have:

$$\leq \sum_{l\in[m]} \frac{d_{il}^{(t)}}{p_l^{(t)}}(\Delta p_l)^2 + \frac{5\epsilon}{12}\frac{4}{3}\sum_{l\in[m]} \frac{d_{il}^{(t)}}{p_l^{(t)}}(\Delta p_l)^2 + \frac{25\epsilon^2}{576}\frac{4}{3}\sum_{l\in[m]} \frac{d_{il}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)}{p_l^{(t)}}(\Delta p_l)^2 \tag{7.32}$$

$$= \left(1 + \frac{5\epsilon}{9}\right)\sum_{l\in[m]} \frac{d_{il}^{(t)}}{p_l^{(t)}}(\Delta p_l)^2 + \left(\frac{25\epsilon^2}{432}\right)\sum_{l\in[m]} \frac{d_{il}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i)}{p_l^{(t)}}(\Delta p_l)^2 \tag{7.33}$$

Lemma 7.2.12.
Suppose that $\frac{|\Delta p_j|}{p_j^{(t)}} \leq \frac{1}{4}$, then

$$b_i \log\left(1 - \frac{e_i(\boldsymbol{p}^{(t+1)}, 1) - e_i(\boldsymbol{p}^{(t)}, 1)}{e_i(\boldsymbol{p}^{(t)}, 1)}\left(1 + \frac{e_i(\boldsymbol{p}^{(t+1)}, 1) - e_i(\boldsymbol{p}^{(t)}, 1)}{e_i(\boldsymbol{p}^{(t)}, 1)}\right)^{-1}\right)$$
$$\leq \left(\frac{4}{3} + \frac{20\epsilon}{27}\right)\sum_{l\in[m]} \frac{d_{il}^{(t)}}{p_l^{(t)}}(\Delta p_l)^2 + \left(\frac{5\epsilon}{6} + \frac{25\epsilon^2}{324}\right)\sum_{l\in[m]} \frac{(\Delta p_l)^2}{p_l^{(t)}} d_{il}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i) - \left\langle \boldsymbol{d}_i(\boldsymbol{p}^{(t)}, 1), \Delta \boldsymbol{p}\right\rangle$$

Proof of Lemma 7.2.12

First, we note that $\boldsymbol{h}_i^{(t)} \cdot \boldsymbol{p}^{(t)} > 0$ because prices during our tâtonnement rule reach 0 only asymptotically and Hicksian demand for one unit of utility at prices $\boldsymbol{p}^{(t)} > 0$ is strictly positive; and likewise, prices reach ∞ only asymptotically, which implies that Hicksian demand is always strictly positive. This fact will come handy, as we divide some expressions by $\boldsymbol{h}_i^{(t)} \cdot \boldsymbol{p}^{(t)}$.

Fix $t \in \mathbb{N}_+$ and $i \in [n]$. Since by our assumptions $\frac{|\Delta p_j|}{p_j^{(t)}} \leq \frac{1}{4}$, by Lemma 7.2.10, we have $0 \leq \left|\frac{e_i(\boldsymbol{p}^{(t+1)},1) - e_i(\boldsymbol{p}^{(t)},1)}{e_i(\boldsymbol{p}^{(t)},1)}\right| \leq \frac{1}{4}$. We can then use the bound $1 - x(1+x)^{-1} \leq 1 + \frac{4}{3}x^2 - x$, for $0 \leq |x| \leq \frac{1}{4}$, with $x = \frac{e_i(\boldsymbol{p}^{(t+1)},1) - e_i(\boldsymbol{p}^{(t)},1)}{e_i(\boldsymbol{p}^{(t)},1)}$, to get:

$$
\begin{aligned}
&b_i \log\left(1 - \frac{e_i(\boldsymbol{p}^{(t+1)},1) - e_i(\boldsymbol{p}^{(t)},1)}{e_i(\boldsymbol{p}^{(t)},1)}\left(1 + \frac{e_i(\boldsymbol{p}^{(t+1)},1) - e_i(\boldsymbol{p}^{(t)},1)}{e_i(\boldsymbol{p}^{(t)},1)}\right)^{-1}\right) \\
&\leq b_i \log\left(1 + \frac{4}{3}\left(\frac{e_i(\boldsymbol{p}^{(t+1)},1) - e_i(\boldsymbol{p}^{(t)},1)}{e_i(\boldsymbol{p}^{(t)},1)}\right)^2 - \frac{e_i(\boldsymbol{p}^{(t+1)},1) - e_i(\boldsymbol{p}^{(t)},1)}{e_i(\boldsymbol{p}^{(t)},1)}\right)
\end{aligned}
$$

Let $a = \frac{4}{3}\left(\frac{e_i(\boldsymbol{p}^{(t+1)},1) - e_i(\boldsymbol{p}^{(t)},1)}{e_i(\boldsymbol{p}^{(t)},1)}\right)^2 - \frac{e_i(\boldsymbol{p}^{(t+1)},1) - e_i(\boldsymbol{p}^{(t)},1)}{e_i(\boldsymbol{p}^{(t)},1)}$. By Lemma 7.2.10, we know that $0 + (-1/4) \leq a \leq \frac{1}{12} + 1/4 \Leftrightarrow -1/4 \leq a \leq \frac{1}{3}$. We now use the bound $x \geq \log(1+x)$ for $x > -1$, with $x = a$ to get:

$$
\begin{aligned}
&b_i \log\left(1 + \frac{4}{3}\left(\frac{e_i(\boldsymbol{p}^{(t+1)},1) - e_i(\boldsymbol{p}^{(t)},1)}{e_i(\boldsymbol{p}^{(t)},1)}\right)^2 - \frac{e_i(\boldsymbol{p}^{(t+1)},1) - e_i(\boldsymbol{p}^{(t)},1)}{e_i(\boldsymbol{p}^{(t)},1)}\right) \\
&\leq b_i \left(\frac{4}{3}\left(\frac{e_i(\boldsymbol{p}^{(t+1)},1) - e_i(\boldsymbol{p}^{(t)},1)}{e_i(\boldsymbol{p}^{(t)},1)}\right)^2 - \frac{e_i(\boldsymbol{p}^{(t+1)},1) - e_i(\boldsymbol{p}^{(t)},1)}{e_i(\boldsymbol{p}^{(t)},1)}\right)
\end{aligned}
$$

Using a first order Taylor expansion of $e_i(\boldsymbol{p}^{(t)} + \Delta\boldsymbol{p}, 1)$ around $\boldsymbol{p}^{(t)}$, by Taylor's theorem (Graves, 1927), we have: $e_i(\boldsymbol{p}^{(t)} + \Delta\boldsymbol{p}, 1) = e_i(\boldsymbol{p}^{(t)}, 1) + \left\langle \nabla_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)}, 1), \Delta\boldsymbol{p}\right\rangle + 1/2\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle$ for some $c \in (0,1)$. Re-organizing terms around, we get $e_i(\boldsymbol{p}^{(t+1)}, 1) - e_i(\boldsymbol{p}^{(t)}, 1) = \left\langle \nabla_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)}, 1), \Delta\boldsymbol{p}\right\rangle + 1/2\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle$, which gives us:

$$
\begin{aligned}
&= b_i\left(\frac{4}{3}\left(\frac{\left\langle \nabla_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)}, 1), \Delta\boldsymbol{p}\right\rangle + 1/2\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)}\right)^2 - \right. \\
&\left.\frac{\left\langle \nabla_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)}, 1), \Delta\boldsymbol{p}\right\rangle + 1/2\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p}\right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)}\right) \qquad (7.34)
\end{aligned}
$$

Continuing, by Shepherd's lemma (Shephard, 2015), we have:

$$= \frac{4}{3} b_i \left(\frac{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \Delta \boldsymbol{p} \right\rangle + {}^1\!/_2 \left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)} \right)^2 -$$

$$b_i \frac{\left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \Delta \boldsymbol{p} \right\rangle + {}^1\!/_2 \left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)} \tag{7.35}$$

$$= \frac{4}{3} \frac{1}{b_i} \left(\frac{b_i \left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \Delta \boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)} + \frac{{}^{b_i}\!/_2 \left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)} \right)^2 -$$

$$\frac{b_i \left\langle \boldsymbol{h}_i(\boldsymbol{p}^{(t)}, 1), \Delta \boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)} - \frac{{}^{b_i}\!/_2 \left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)} \tag{7.36}$$

$$= \frac{4}{3} \frac{1}{b_i} \left(\left\langle \boldsymbol{d}_i(\boldsymbol{p}^{(t)}, 1), \Delta \boldsymbol{p} \right\rangle + \frac{{}^{b_i}\!/_2 \left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)} \right)^2 -$$

$$\left\langle \boldsymbol{d}_i(\boldsymbol{p}^{(t)}, 1), \Delta \boldsymbol{p} \right\rangle - \frac{{}^{b_i}\!/_2 \left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)} \tag{7.37}$$

where the last line was obtained from Lemma 7.2.9.

Using Lemma 7.2.11, we have:

$$\leq \frac{4}{3} \left(\left(1 + \frac{5\epsilon}{9}\right) \sum_{l \in [m]} \frac{d_{il}^{(t)}}{p_l^{(t)}} (\Delta p_l)^2 + \frac{25\epsilon^2}{432} \sum_{l \in [m]} \frac{(\Delta p_l)^2}{p_l^{(t)}} d_{il}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, b_i) \right) -$$

$$\left\langle \boldsymbol{d}_i(\boldsymbol{p}^{(t)}, 1), \Delta \boldsymbol{p} \right\rangle - \frac{{}^{b_i}\!/_2 \left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)} \tag{7.38}$$

$$= \left(\frac{4}{3} + \frac{20\epsilon}{27}\right) \sum_{l \in [m]} \frac{d_{il}^{(t)}}{p_l^{(t)}} (\Delta p_l)^2 + \frac{25\epsilon^2}{324} \sum_{l \in [m]} \frac{(\Delta p_l)^2}{p_l^{(t)}} d_{il}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, b_i) -$$

$$\left\langle \boldsymbol{d}_i(\boldsymbol{p}^{(t)}, 1), \Delta \boldsymbol{p} \right\rangle - \frac{{}^{b_i}\!/_2 \left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)} \tag{7.39}$$

Finally, we note that $\nabla^2_{\boldsymbol{p}} e_i$ is negative semi-definite, meaning that we have $\left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle \leq 0$, allowing us to re-express Equation (7.39) as follows:

$$= \left(\frac{4}{3} + \frac{20\epsilon}{27}\right) \sum_{l \in [m]} \frac{d_{il}^{(t)}}{p_l^{(t)}} (\Delta p_l)^2 + \frac{25\epsilon^2}{324} \sum_{l \in [m]} \frac{(\Delta p_l)^2}{p_l^{(t)}} d_{il}(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, b_i) -$$

$$\left\langle \boldsymbol{d}_i(\boldsymbol{p}^{(t)}, 1), \Delta \boldsymbol{p} \right\rangle + \left| \frac{{}^{b_i}\!/_2 \left\langle \nabla^2_{\boldsymbol{p}} e_i(\boldsymbol{p}^{(t)} + c\Delta\boldsymbol{p}, 1)\Delta\boldsymbol{p}, \Delta\boldsymbol{p} \right\rangle}{e_i(\boldsymbol{p}^{(t)}, 1)} \right| \tag{7.40}$$

$$\leq \left(\frac{4}{3} + \frac{20\epsilon}{27}\right) \sum_{l \in [m]} \frac{d_{il}^{(t)}}{p_l^{(t)}} (\Delta p_l)^2 + \frac{25\epsilon^2}{324} \sum_{l \in [m]} \frac{(\Delta p_l)^2}{p_l^{(t)}} d_{il}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i) - \left\langle \boldsymbol{d}_i(\boldsymbol{p}^{(t)}, 1), \Delta \boldsymbol{p} \right\rangle +$$
$$\frac{5\epsilon}{6} \sum_{l \in [m]} \frac{(\Delta p_l)^2}{p_l} d_{il}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i) \tag{7.41}$$
$$\leq \left(\frac{4}{3} + \frac{20\epsilon}{27}\right) \sum_{l \in [m]} \frac{d_{il}^{(t)}}{p_l^{(t)}} (\Delta p_l)^2 +$$
$$\left(\frac{5\epsilon}{6} + \frac{25\epsilon^2}{324}\right) \sum_{l \in [m]} \frac{(\Delta p_l)^2}{p_l^{(t)}} d_{il}(\boldsymbol{p}^{(t)} + c\Delta \boldsymbol{p}, b_i) - \left\langle \boldsymbol{d}_i(\boldsymbol{p}^{(t)}, 1), \Delta \boldsymbol{p} \right\rangle \tag{7.42}$$

where the penultimate line was obtained from Lemma 6.4.2.

Lemma 6.5.2 [Bounded Indirect Utility for Homothetic Fisher Markets]**.**

If entropic *tâtonnement* is run on a homothetic Fisher market $(\boldsymbol{u}, \boldsymbol{b})$, then, for all $t \in \mathbb{N}_+$, the following bound holds:

$$v_i(\boldsymbol{p}^{(t)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) \leq v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) + 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{\frac{h_{ij}(\boldsymbol{q}, 1)^2}{h_{ik}(\boldsymbol{p}, 1)^2}\right\}$$

Proof of Lemma 6.5.2

Fix a buyer $i \in [n]$. First, note that since by Lemma 7.2.1, since the expenditure function is homogeneous of degree 0 in prices, we have for all $j \in [m]$, $\max_{\boldsymbol{p} \in \mathbb{R}_+^m} h_{ij}(\boldsymbol{p}, 1) = \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1)$. In addition, note that $\max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1)$ is well-defined since Δ_m is compact, $h_{ij}(\boldsymbol{p}, 1)$ exists for all $\boldsymbol{p} \in \mathbb{R}_+^m$, and is by Berge's maximum theorem (Berge, 1997) continuous in homothetic Fisher markets.

We will now prove that for any $t \in \mathbb{N}$, $v_i(\boldsymbol{p}^{(t)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) \leq v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) + 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{\frac{h_{ij}(\boldsymbol{q}, 1)^2}{h_{ik}(\boldsymbol{p}, 1)^2}\right\}$ by induction.

Base case: $t = 0$**.** By definition, we have

$$v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) \leq v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) + 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{\frac{h_{ij}(\boldsymbol{q}, 1)^2}{h_{ik}(\boldsymbol{p}, 1)^2}\right\} .$$

Inductive hypothesis. Suppose that for any $t \in \mathbb{N}$, we have:

$$v_i(\boldsymbol{p}^{(t)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) \leq v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{p} \in \Delta_m} h_{ij}(\boldsymbol{p}, 1) + 2 \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m]: h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{\frac{h_{ij}(\boldsymbol{q}, 1)^2}{h_{ik}(\boldsymbol{p}, 1)^2}\right\}$$

Inductive step. We will show that the inductive hypothesis holds for $t+1$. We proceed with a proof by cases.

Case 1: $d_{ij}^{(t)} \geq \max_{\substack{\boldsymbol{p},\boldsymbol{q}\in\Delta_m \\ k\in[m]:h_{ik}(\boldsymbol{p},1)>0}} \left\{\frac{h_{ij}(\boldsymbol{q},1)}{h_{ik}(\boldsymbol{p},1)}\right\}$.

For all $k \in [m]$, we have:

$$\begin{aligned}
d_{ik}^{(t)} &= \frac{h_{ik}^{(t)}}{h_{ij}^{(t)}} d_{ij}^{(t)} \\
&\geq \frac{h_{ik}^{(t)}}{h_{ij}^{(t)}} \max_{\substack{\boldsymbol{p},\boldsymbol{q}\in\Delta_m \\ k\in[m]:h_{ik}(\boldsymbol{p},1)>0}} \left\{\frac{h_{ij}(\boldsymbol{q},1)}{h_{ik}(\boldsymbol{p},1)}\right\} \\
&\geq \frac{h_{ik}^{(t)}}{h_{ij}^{(t)}} \frac{h_{ij}^{(t)}}{h_{ik}^{(t)}} = 1
\end{aligned}$$

where the penultimate line follows from the case hypothesis.

The above means that the price of all goods will increase in the next time period, i.e., $\forall k \in [m], p_k^{(t+1)} \geq p_k^{(t)}$ which implies that $e_i(\boldsymbol{p}^{(t+1)}, 1) \geq e_i(\boldsymbol{p}^{(t)}, 1) \geq 0$. In addition, note that the expenditure is positive since prices reach 0 only asymptotically under entropic *tâtonnement*. Which gives us:

$$\begin{aligned}
\frac{b_i}{e_i(\boldsymbol{p}^{(t+1)}, 1)} &\leq \frac{b_i}{e_i(\boldsymbol{p}^{(t)}, 1)} \\
v_i(\boldsymbol{p}^{(t+1)}, b_i) &\leq v_i(\boldsymbol{p}^{(t)}, b_i) && \text{(Corollary 6.2.1)}
\end{aligned}$$

Multiplying both sides by $\max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1)$, we have for all $j \in [m]$:

$$\begin{aligned}
v_i(\boldsymbol{p}^{(t+1)}, b_i) \max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1) &\leq v_i(\boldsymbol{p}^{(t)}, b_i) \max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1) \\
&= v_i(\boldsymbol{p}^{(0)}, b_i) \max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p}, 1) + 2 \max_{\substack{\boldsymbol{p},\boldsymbol{q}\in\Delta_m \\ k\in[m]:h_{ik}(\boldsymbol{p},1)>0}} \left\{\frac{h_{ij}(\boldsymbol{q},1)^2}{h_{ik}(\boldsymbol{p},1)^2}\right\}
\end{aligned}$$

where the last line follows by the induction hypothesis.

Case 2: $d_{ij}^{(t)} < \max_{\substack{\boldsymbol{p},\boldsymbol{q}\in\Delta_m \\ k\in[m]:h_{ik}(\boldsymbol{p},1)>0}} \left\{\frac{h_{ij}(\boldsymbol{q},1)}{h_{ik}(\boldsymbol{p},1)}\right\}$.

For all $k \in [m]$, we have:

$$
\begin{aligned}
d_{ik}^{(t)} &= \frac{h_{ik}^{(t)}}{h_{ij}^{(t)}} d_{ij}^{(t)} \\
&\leq \frac{h_{ik}^{(t)}}{h_{ij}^{(t)}} \max_{\substack{\boldsymbol{p}, \boldsymbol{q} \in \Delta_m \\ k \in [m] : h_{ik}(\boldsymbol{p}, 1) > 0}} \left\{ \frac{h_{ij}(\boldsymbol{q}, 1)}{h_{ik}(\boldsymbol{p}, 1)} \right\} \\
&= \frac{h_{ik}^{(t)}}{h_{ij}^{(t)}} \frac{h_{ij}^{(t)}}{h_{ik}^{(t)}} = 1
\end{aligned}
$$

where the penultimate line follows from the case hypothesis.

The above means that prices of all goods will decrease in the next time period. Now, note that regardless of the aggregate demand $\boldsymbol{q}^{(t)}$ at time $t \in \mathbb{N}$, prices can decrease at most by a factor of $e^{-\frac{1}{5}} \geq 1/2$, that is, for all $j \in [m]$

$$
\begin{aligned}
p_j^{(t+1)} &= p_j^{(t)} \exp \left\{ \frac{z_j(\boldsymbol{p}^{(t)})}{\gamma} \right\} \\
&= p_j^{(t)} \exp \left\{ \frac{q_j^{(t)} - 1}{\gamma} \right\} \\
&\geq p_j^{(t)} \exp \left\{ \frac{-1}{\gamma} \right\} \\
&\geq p_j^{(t)} \exp \left\{ \frac{-1}{5 \max\limits_{\substack{t \in \mathbb{N} \\ j \in [m]}} \{1, q_j^{(t)}\}} \right\} \\
&\geq p_j^{(t)} \exp \left\{ \frac{-1}{5} \right\} \\
&\geq p_j^{(t)} e^{-\frac{1}{5}} \geq \frac{1}{2} p_j^{(t)}
\end{aligned}
$$

Now, notice that we have $e_i(\boldsymbol{p}^{(t+1)}, 1) \geq e_i(\frac{1}{2}\boldsymbol{p}^{(t)}, 1) = \frac{1}{2} e_i(\boldsymbol{p}^{(t)}, 1) \geq 0$, since the expenditure of the buyer increases the most when the prices of all goods decrease simultaneously and the expenditure function is homogeneous of degree 1 in prices. In addition, note that the expenditure is positive since prices reach 0 only asymptotically under entropic *tâtonnement*. Hence, we have:

$$\frac{b_i}{e_i(\boldsymbol{p}^{(t+1)},1)} \leq 2\frac{b_i}{e_i(\boldsymbol{p}^{(t)},1)}$$

$$v_i(\boldsymbol{p}^{(t+1)},b_i) \leq 2v_i(\boldsymbol{p}^{(t)},b_i) \qquad \text{(Corollary 6.2.1)}$$

Multiplying both sides by $h_{ij}^{(t)}$, and applying Lemma 7.2.9, we have for all $j \in [m]$:

$$v_i(\boldsymbol{p}^{(t+1)},b_i)h_{ij}^{(t)} \leq 2d_{ij}^{(t)}$$

$$v_i(\boldsymbol{p}^{(t+1)})h_{ij}^{(t)} \leq 2 \max_{\substack{\boldsymbol{p},\boldsymbol{q}\in\Delta_m \\ k\in[m]:h_{ik}(\boldsymbol{p},1)>0}} \left\{\frac{h_{ij}(\boldsymbol{q},1)}{h_{ik}(\boldsymbol{p},1)}\right\} \qquad \text{(Case hypothesis)}$$

Now, taking a maximum over all $j \in [m]$ s.t. $h_{ij}^{(t)} > 0$, we have

$$v_i(\boldsymbol{p}^{(t+1)}) \max_{k\in[m]:h_{ik}^{(t)}>0} h_{ik}^{(t)} \leq 2 \max_{\substack{\boldsymbol{p},\boldsymbol{q}\in\Delta_m \\ k\in[m]:h_{ik}(\boldsymbol{p},1)>0}} \left\{\frac{h_{ij}(\boldsymbol{q},1)}{h_{ik}(\boldsymbol{p},1)}\right\}$$

$$v_i(\boldsymbol{p}^{(t+1)}) \max_{\substack{\boldsymbol{p}\in\Delta_m \\ k\in[m]:h_{ik}(\boldsymbol{p},1)>0}} h_{ik}(\boldsymbol{p},1) \leq 2 \max_{\substack{\boldsymbol{p},\boldsymbol{q}\in\Delta_m \\ k\in[m]:h_{ik}(\boldsymbol{p},1)>0}} \left\{\frac{h_{ij}(\boldsymbol{q},1)}{h_{ik}(\boldsymbol{p},1)}\right\}$$

$$v_i(\boldsymbol{p}^{(t+1)}) \leq \frac{2}{\max_{\substack{\boldsymbol{p}\in\Delta_m \\ k\in[m]:h_{ik}(\boldsymbol{p},1)>0}} h_{ik}(\boldsymbol{p},1)} \max_{\boldsymbol{p},\boldsymbol{q}} \max_{k\in[m]:h_{ik}(\boldsymbol{p},1)>0} \left\{\frac{h_{ij}(\boldsymbol{q},1)}{h_{ik}(\boldsymbol{p},1)}\right\}$$

Finally, multiplying both sides by $\max_{\boldsymbol{q}\in\Delta_m} h_{ij}(\boldsymbol{q},1)$, we have:

$$v_i(\boldsymbol{p}^{(t+1)}) \max_{\boldsymbol{q}\in\Delta_m} h_{ij}(\boldsymbol{q},1) \leq 2\frac{\max_{\boldsymbol{q}\in\Delta_m} h_{ij}(\boldsymbol{q},1)}{\max_{\substack{\boldsymbol{p}\in\Delta_m \\ k\in[m]:h_{ik}(\boldsymbol{p},1)>0}} h_{ik}(\boldsymbol{p},1)} \max_{\boldsymbol{p},\boldsymbol{q}} \max_{k\in[m]:h_{ik}(\boldsymbol{p},1)>0} \left\{\frac{h_{ij}(\boldsymbol{q},1)}{h_{ik}(\boldsymbol{p},1)}\right\}$$

$$v_i(\boldsymbol{p}^{(t+1)}) \max_{\boldsymbol{q}\in\Delta_m} h_{ij}(\boldsymbol{q},1) \leq 2 \max_{\substack{\boldsymbol{p},\boldsymbol{q}\in\Delta_m \\ k\in[m]:h_{ik}(\boldsymbol{p},1)>0}} \left\{\frac{h_{ij}(\boldsymbol{q},1)}{h_{ik}(\boldsymbol{p},1)}\right\}^2$$

$$v_i(\boldsymbol{p}^{(t+1)}) \max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p},1) \leq v_i(\boldsymbol{p}^{(0)},b_i) \max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p},1) + 2 \max_{\substack{\boldsymbol{p},\boldsymbol{q}\in\Delta_m \\ k\in[m]:h_{ik}(\boldsymbol{p},1)>0}} \left\{\frac{h_{ij}(\boldsymbol{q},1)^2}{h_{ik}(\boldsymbol{p},1)^2}\right\}$$

Hence, the inductive hypothesis holds for $t+1$. Putting it all together, we have, for all $t \in \mathbb{N}$:

$$\sum_{i\in[n]} v_i(\boldsymbol{p}^{(t)},b_i) \max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p},1) \leq \sum_{i\in[n]} v_i(\boldsymbol{p}^{(0)},b_i) \max_{\boldsymbol{p}\in\Delta_m} h_{ij}(\boldsymbol{p},1) + 2\sum_{i\in[n]} \max_{\substack{\boldsymbol{p},\boldsymbol{q}\in\Delta_m \\ k\in[m]:h_{ik}(\boldsymbol{p},1)>0}} \left\{\frac{h_{ij}(\boldsymbol{q},1)^2}{h_{ik}(\boldsymbol{p},1)^2}\right\}$$

Part II

Pseudo-Games and Arrow-Debreu Exchange Economies

Chapter 8

Scope and Motivation

8.1 Scope

Part II of this thesis is divided into two chapters. In Chapter 9, after reviewing background material on pseudo-games, I will introduce three new algorithmic approaches for computing equilibria in pseudo-games with polynomial-time guarantees. The first approach consists of a family of first-order algorithms known as mirror extragradient learning dynamics. I will prove that these methods converge to a variational equilibrium (VE) in variationally stable concave pseudo-games with jointly convex constraints. Beyond concave settings, I will establish convergence to a first-order variational equilibrium. Next, I will introduce two types of merit function minimization methods—one first-order and one second-order—that compute a solution satisfying the necessary conditions for a variational equilibrium in Lipschitz-smooth pseudo-games with jointly convex constraints.

In Chapter 10, after reviewing the foundational model of Arrow-Debreu economies, I will demonstrate that the set of Arrow-Debreu equilibria in any pure exchange economy corresponds exactly to the set of generalized Nash equilibria (GNE) of an associated variationally stable pseudo-game with jointly convex constraints. Leveraging this equivalence, I will introduce a novel family of market dynamics, called mirror extratrade dynamics, and prove their polynomial-time convergence to an Arrow-Debreu equilibrium in pure exchange economies. Finally, for more general, possibly non-concave, Arrow-Debreu economies, I will develop two polynomial-time merit function methods that compute a solution satisfying the necessary conditions for an Arrow-Debreu equilibrium.

8.2 Motivation

Walrasian equilibrium (Arrow and Debreu, 1954; Walras, 1896), first studied by French economist Léon Walras in 1874, is a steady state of an economy—any system governed by supply and demand (Walras, 1896). Walras assumed that, given prices, each producer in an economy would act so as to maximize its profit, while consumers would make decisions that maximize their preferences over their available consumption choices; all this, while perfect competition prevails, meaning producers and consumers are unable to influence the prices that emerge. Under these assumptions, the demand and supply of each commodity is a function of prices, as they are a consequence of the decisions made by the producers and consumers, having observed the prevailing prices. A Walrasian equilibrium then corresponds to prices that solve the system of simultaneous equations with demand on one side and supply on the other, i.e., prices at which supply meets demand. Unfortunately, Walras did not provide conditions that guarantee the existence of such a solution, and the question of whether such prices exist remained open until Arrow and Debreu's rigorous analysis of Walrasian equilibrium in their model of a competitive economy in the middle of last century (Arrow and Debreu, 1954).

The Arrow-Debreu model comprises a set of commodities; a set of firms, each deciding what quantity of each commodity to supply; and a set of consumers, each choosing a quantity of each commodity to demand in exchange for their endowment (Arrow and Debreu, 1954). Arrow and Debreu define an **Arrow-Debreu equilibrium** as a collection of consumptions, one per consumer, a collection of productions, one per firm, and a collection of prices, one per commodity, such that fixing equilibrium prices: (1) no consumer can increase their utility by deviating to an alternative affordable consumption, (2) no firm can increase profit by deviating to another production in their production set, and (3) the **aggregate demand** for each commodity (i.e., the sum of the commodity's consumption across all consumers) does not exceed to its **aggregate supply** (i.e., the sum of the commodity's production and endowment across firms and consumers, respectively), while the total value of the aggregate demand is equal to the total value of the aggregate supply, i.e., **Walras' law** holds. That is, an Arrow-Debreu equilibrium is a tuple that comprises a Walrasian equilibrium, and the associated utility-maximizing consumptions and profit-maximizing productions.

Arrow and Debreu proceeded to show that their competitive economy could be seen as an **abstract economy**, which today is better known as a **pseudo-game** (Arrow and Debreu, 1954; Facchinei and Kanzow, 2010a). A pseudo-game is a generalization of a game in which the actions taken by each player impact not only the other players' payoffs, as in games, but also their set of permissible actions. Pseudo-games generalize games, and hence are even more widely applicable. Some recently studied applications include adversarial classification (Bruckner et al., 2012; Bruckner and Scheffer, 2009), energy resource allocation (Hobbs and Pang, 2007; Jing-Yuan and Smeers, 1999), environmental protection (Breton et al., 2006; Krawczyk, 2005), cloud computing (Ardagna et al., 2017; 2011), ride sharing services ((Jeff) Ban et al., 2019), transportation (Stein and Sudermann-Merx, 2018), and wireless and network communication (Han et al., 2011; Pang et al., 2008).

Arrow and Debreu proposed **generalized Nash equilibrium** as the solution concept for this model, an action profile from which no player can improve their payoff by unilaterally deviating to another action in the space of permissible actions determined by the actions of other players. Arrow and Debreu further showed that any competitive economy could be represented as a pseudo-game inhabited by a fictional auctioneer, who sets prices so as to buy and resell commodities at a profit, as well as consumers and producers, who respectively, choose utility-maximizing consumptions of commodities in the budget sets determined by the prices set by the auctioneer, and profit-maximizing productions at the prices set by the auctioneer. The elegance of the reduction from competitive economies to pseudo-games is rooted in a simple observation: the set of Arrow-Debreu equilibria of a competitive economy is equal to the set of generalized Nash equilibria of the associated pseudo-game, implying the existence of Arrow-Debreu equilibrium in competitive economies (and hence, Walrasian equilibrium in Walrasian economies as a corollary of the existence of generalized Nash equilibria in pseudo-games, whose proof is a straightforward generalization of Nash's proof for the existence of Nash equilibria (Nash, 1950b).[1]

Following Arrow and Debreu's seminal existence result, the literature turned its attention to questions of 1. (economic) **efficiency**, i.e., under what assumptions are Arrow-Debreu equilibria Pareto-optimal? (Arrow, 1951a;b; Arrow and Nerlove, 1958; Arrow and Hurwicz, 1958; Balasko, 1975; Debreu, 1951a); 2. **uniqueness**, i.e., under what assumptions are Arrow-Debreu equilibria unique? (Dierker, 1982; Pearce and Wise, 1973); 3. **stability**, i.e., under what conditions can a competitive economy settle into an Arrow-Debreu equilibrium? (Hahn, 1958; Balasko, 1975; Arrow and Hurwicz, 1958; Cole and Fleischer, 2008; Cheung et al., 2018; 2013; Goktas et al., 2023c), and 4. **efficient computation**, i.e., under what conditions can an Arrow-Debreu equilibrium be computed efficiently? (Jain et al., 2005; Codenotti et al., 2005; 2006; Chen and Teng, 2009).

In this part of the thesis, we seek to provide an answer to the latter two questions by introducing a family of algorithms to compute a generalized Nash equilibrium in pseudo-games. Work in this direction is progressing; see, for example, (Facchinei et al., 2009; Facchinei and Kanzow, 2010a; Facchinei and Sagratella, 2011; Paccagnan et al., 2016; Yi and Pavel, 2017; Couzoudis and Renner, 2013; Dreves, 2017; Von Heusinger and Kanzow, 2009; Tatarenko and Kamgarpour, 2018; Dreves and Sudermann-Merx, 2016; Von Heusinger et al., 2012; Izmailov and Solodov, 2014; Fischer et al., 2016; Pang and Fukushima, 2005; Facchinei and Lampariello, 2011; Fukushima, 2011; Kanzow, 2016; Kanzow and Steck, 2016; 2018; Ba and Pang, 2020). Nonetheless, there are still few, if any (Jordan et al., 2022), GNE-finding algorithms with computational guarantees, even for restricted classes of pseudo-games.

[1]McKenzie (1959) showed the existence of Walrasian equilibrium independently, but concurrently. Much of his work, however, has gone unrecognized perhaps because his proof technique does not depend on this fundamental relationship between competitive and abstract economies / pseudo-games.

8.3 Contributions

8.3.1 Pseudo-Games

In Chapter 9, I advance the study of pseudo-games by refining their solution concepts and analyzing their computational complexity. First, I re-establish the existence of variational equilibrium in quasiconcave pseudo-games with jointly convex constraints (Theorem 9.2.2). I then introduce the notion of first-order variational equilibrium, which I show exists in a broader class of pseudo-games—namely, smooth games with jointly convex constraints (Theorem 9.5.2). Next, I establish an equivalence between (first-order) variational equilibria of pseudo-games and strong solutions of variational inequalities (Lemma 9.4.1 and Lemma 9.6.1). This allows me to define a new class of pseudo-games—variationally stable pseudo-games with jointly convex constraints—for which a first-order variational equilibrium can be computed in polynomial time via a novel uncoupled learning dynamic called the mirror extragradient learning dynamics (Theorem 9.6.1). In the special case where the pseudo-game is also concave, this result extends to the computation of variational equilibrium in polynomial time via these learning dynamics (Theorem 9.4.1). To the best of my knowledge, this result is the broadest of its kind in the literature. Finally, for more general pseudo-games with jointly convex constraints that are not necessarily variationally stable, I develop two polynomial-time globally convergent merit function methods that compute a solution satisfying the necessary conditions for a variational equilibrium (Theorem 9.4.3 and Theorem 9.6.2).

8.3.2 Arrow-Debreu Economies

In Chapter 10, I provide novel mathematical characterizations of Arrow-Debreu equilibrium in Arrow-Debreu economies. First, I re-establish that the set of Arrow-Debreu equilibria of any quasiconcave Arrow-Debreu economy coincides with the set of generalized Nash equilibria of the corresponding Arrow-Debreu pseudo-game (Lemma 10.2.1). However, as the Arrow-Debreu pseudo-game characterization is intractable, I introduce an alternative characterization: the set of Arrow-Debreu equilibria of any concave pure exchange economy corresponds to the set of generalized Nash equilibria of the trading post pseudo-game (Lemma 10.3.2), which is a variationally stable pseudoconcave pseudo-game with jointly convex constraints. I then apply the mirror extragradient learning dynamics to solve this pseudo-game, leading to a market dynamic I call mirror extratrade dynamics. While the trading post pseudo-game is not concave, I show that it is pseudoconcave, implying that an approximate first-order variational equilibrium can be computed in polynomial time. Moreover, asymptotically, the algorithm converges to a variational equilibrium of the trading post pseudo-game—and thus to an Arrow-Debreu equilibrium of the associated concave pure exchange economy (Theorem 10.3.1). Finally, for more general, possibly non-concave, Arrow-Debreu economies, I develop two polynomial-time globally convergent merit function methods that compute a solution satisfying the necessary conditions for an Arrow-Debreu equilibrium (Theorem 10.4.1 and Theorem 10.4.2).

My results on the computation of Arrow-Debreu equilibrium do not contradict any known PPAD-hardness results about Arrow-Debreu economies. The reason for this is, the approximate equilibria of the trading post pseudo-game do not coincide with the approximate competitive equilibria of an Arrow-Debreu economy. In other words, under mild continuity and smoothness assumptions, I have established polynomial-time convergence to an *alternative* solution concept in Arrow-Debreu economies; yet, this alternative solution concept is meaningful, in the sense that it is encountered en route to computing an exact Arrow-Debreu equilibrium, which the method finds asymptotically! This result supports a recent research direction proposed by Costis Daskalakis (Daskalakis, 2022), where he advocates for the design of meaningful polynomial-time computable solutions to non-concave games.

Chapter 9

Pseudo-games

9.1 Background

A **pseudo-game** (Arrow and Debreu, 1954) $(n, l, \mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$, denoted $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ when n and l are clear from context, comprises $n \in \mathbb{N}_+$ players, each player $i \in [n]$ of which takes an **action** $\boldsymbol{a}_i \in \mathcal{A}_i$ from its **action space** $\mathcal{A}_i$. An ordered tuple of per-player actions $\boldsymbol{a} \doteq (\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n) \in \mathcal{A}$ is called an **action profile**, where we define $\mathcal{A} \doteq \bigtimes_{i \in [n]} \mathcal{A}_i \subset \mathbb{R}^{nm}$ to be the space of action profiles. We denote any action profile $\boldsymbol{a} \in \mathcal{A}$ where the ith player's action is removed by $\boldsymbol{a}_{-i} \in \mathcal{A}_{-i}$, where $\mathcal{A}_{-i} \doteq \bigtimes_{i' \neq i} \mathcal{A}_{i'} \subset \mathbb{R}^{(n-1)m}$. Additionally, we use the notation $(\boldsymbol{a}'_i, \boldsymbol{a}_{-i}) \in \mathcal{A}$ to denote the action profile $\boldsymbol{a} \in \mathcal{A}$, where the ith player's action is replaced by the action $\boldsymbol{a}_i \in \mathcal{A}_i$.

Each player i simultaneously chooses a **feasible** action from the set $\mathcal{X}_i(\boldsymbol{a}_{-i}) = \{\boldsymbol{a}_i \in \mathcal{A}_i \mid g_{ic}(\boldsymbol{a}_i, \boldsymbol{a}_{-i}) \geq 0, \text{ for all } c \in [l]\}$, determined by the actions $\boldsymbol{a}_{-i} \in \mathcal{A}_{-i}$ of the other players, where $g_{ic} : \mathcal{A} \to \mathbb{R}^l$ is the **action constraint function**, l is the number of constraints, and $\mathcal{X}_i : \mathcal{A}_{-i} \rightrightarrows \mathcal{A}_i$ is the **feasible action correspondence**. We denote the product of these feasible action correspondences by $\mathcal{X}(\boldsymbol{a}) = \bigtimes_{i \in [n]} \mathcal{X}_i(\boldsymbol{a}_{-i})$, and define the **space of feasible action profiles** $\mathcal{X}^* \doteq \{\boldsymbol{a} \in \mathcal{A} \mid \boldsymbol{a} \in \mathcal{X}(\boldsymbol{a})\}$. Once players have taken a feasible action $\boldsymbol{a} \in \mathcal{X}(\boldsymbol{a})$, each player i receives a payoff $u_i(\boldsymbol{a})$ according to their **payoff function** $u_i : \mathcal{A} \to \mathbb{R}$. The **payoff profile function** $\boldsymbol{u}(\boldsymbol{a}) = (u_i(\boldsymbol{a}))_{i \in [n]}$.

9.2 Global Solution Concepts and Existence

9.2.1 Generalized Nash Equilibrium and Variational Equilibrium

The canonical solution concept for a pseudo-game is the generalized Nash equilibrium (GNE).

Definition 9.2.1 [Generalized Nash equilibrium]**.**

Given $\varepsilon \geq 0$, an ε-**generalized Nash equilibrium (GNE)** is an action profile $\boldsymbol{a}^* \in \mathcal{X}(\boldsymbol{a}^*)$ s.t. for all $i \in [n]$ and $\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{a}^*_{-i})$:

$$u_i(\boldsymbol{a}^*) \geq u_i(\boldsymbol{a}_i, \boldsymbol{a}^*_{-i}) - \varepsilon \ .$$

A 0-GNE is simply called a generalized Nash equilibrium (GNE).

The GNE computation problem can succinctly be written as solving the following n simultaneous quasi-optimization[1] problems:

$$\forall i \in [n], \qquad \max_{\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{a}_{-i})} u_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i})$$

An important refinement of GNE is the variational equilibrium (VE):

Definition 9.2.2 [Variational equilibrium]**.**

Given $\varepsilon \geq 0$, a ε-**variational equilibrium (VE)** is an action profile $\boldsymbol{a}^* \in \mathcal{X}(\boldsymbol{a}^*)$ s.t. for all action profiles $\boldsymbol{a} \in \mathcal{X}^*$:

$$\sum_{i \in [n]} \left[u_i(\boldsymbol{a}_i, \boldsymbol{a}^*_{-i}) - u_i(\boldsymbol{a}^*)\right] \leq \varepsilon \ .$$

A 0-VE is simply called a variational equilibrium (VE).

An important class of pseudo-games are those which are unconstrained, more commonly known as games.

Definition 9.2.3 [Games and Nash equilibrium]**.**

A **game** (Nash, 1950b) $(n, \boldsymbol{u}, \mathcal{A})$, denoted $(\boldsymbol{u}, \mathcal{A})$ when n is clear from context, is a pseudo-

[1] A quasi-optimization problem is a computational problem which consists of finding a solution which is 1) a fixed point of a constraint correspondence, 2) a maximizer of the objective over the set of all feasible variables defined by the constraint correspondence evaluated at the solution.

game $(n, l, \mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$, where $l \doteq 0$ and $\boldsymbol{g} = \emptyset$, i.e., the pseudo-game is unconstrained and there is no constraint function.

An ε-generalized Nash equilibrium of a game is simply called an ε-**Nash equilibrium** (**NE**). A 0-Nash equilibrium is simply called a Nash equilibrium.

Similar to the GNE computation problem, the NE computation problem for a game $(\mathcal{A}, \boldsymbol{u})$ can be succinctly expressed the following simultaneous optimization problem:

$$\forall i \in [n], \qquad \max_{\boldsymbol{a}_i \in \mathcal{A}_i} u_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i})$$

Remark 9.2.1 [Relationship between GNE, VE and NE]**.**
While in general, the set of VE is a subset of the set of GNE in pseudo-games, the set of GNE is equal to the set of VE in games, since $\mathcal{X}^* = \mathcal{A}$. The algorithms we provide in this thesis are relevant to VEs in pseudo-games; however, due to the aforementioned relationship between solution concepts, they imply computational results for GNE in pseudo-games where $\mathcal{X}^* = \mathcal{A}$ and for NE in games.

9.2.2 Quasiconcave Pseudo-Games

With a definition of our solution concepts in hand, we now describe the classes of games in which they are guaranteed to exist. The canonical class of pseudo-games in which a GNE is guaranteed to exist is the class of quasiconcave games. Related and of greater interest in the computational literature is the class of concave pseudo-games, which is a strict subset of the class of quasiconcave pseudo-games.

Definition 9.2.4 [Quasiconcave pseudo-games]**.**
A **quasiconcave** (respectively, **concave**) **pseudo-game** is a pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$, where for all players $i \in [n]$:

[Continuous payoffs]	u_i is continuous
[(Quasi)concave payoffs]	$\boldsymbol{a}_i \mapsto u_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i})$ is quasiconcave (respectively, concave) for all $\boldsymbol{a}_{-i} \in \mathcal{A}_{-i}$

[Convex constraints]	$\mathcal{X}_{-i}$ is continuous, non-empty-, compact-, and convex-valued
[Convex action space]	$\mathcal{A}_i$ is non-empty, compact, and convex

Remark 9.2.2 [Convex constraints]**.**
The third condition in the above definition (i.e., the convex constraints condition) can be expressed in terms of the action constraint function $\boldsymbol{g}$, assuming $\boldsymbol{g}$ is continuous and satisfies Slater's condition (as defined in Definition 9.2.6); see the end of Section 2.9 for details.

The importance of concave pseudo-games is due to a seminal result of Arrow and Debreu (1954), which established that a GNE is guaranteed to exist in all quasiconcave pseudo-games. This proof of existence relies on a fixed-point argument, which we provide a short version of here for completeness.

Theorem 9.2.1 [Lemma on abstract economies (Lemma 2.5 of Arrow and Debreu (1954)]**.**
A GNE is guaranteed to exist in all quasiconcave pseudo-games.

Proof of Theorem 9.2.1

Define the **best-response correspondence** of the pseudo-game as: $\mathcal{BR}(\boldsymbol{a}) \doteq \bigtimes_{i \in [n]} \arg\max_{\boldsymbol{a}'_i \in \mathcal{X}_i(\boldsymbol{a}_{-i})} u_i(\boldsymbol{a}'_i, \boldsymbol{a}_{-i})$. Now, consider any fixed point $\boldsymbol{a}^* \in \mathcal{BR}(\boldsymbol{a}^*)$ of the best-response correspondence. Then, for all players $i \in [n]$, we have:

$$\boldsymbol{a}^*_i \in \mathop{\arg\max}_{\boldsymbol{a}'_i \in \mathcal{X}_i(\boldsymbol{a}^*_{-i})} u_i(\boldsymbol{a}'_i, \boldsymbol{a}^*_{-i})$$

That is, $\boldsymbol{a}^*$ is a GNE. Now, by the maximum theorem (Berge, 1997), in quasiconcave games, the best-response correspondence is guaranteed to be upper-hemicontinuous, non-empty-, compact-, and convex-valued. Hence, the best-response correspondence satisfies the assumptions of the Kakutani-Glicksberg fixed point theorem (Theorem 2.4.1), and so a fixed point, and thus a GNE, exists.

Another important class of pseudo-games is the class of pseudo-games with jointly convex constraints.

Definition 9.2.5 [Jointly convex constraints]**.**

A pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is said to have jointly convex constraints iff:

[Non-emptiness] there exists $\boldsymbol{a} \in \mathcal{A}$ s.t. $\boldsymbol{g}(\boldsymbol{a}) \geq \boldsymbol{0}$ (i.e., $\mathcal{X}^*$ is non-empty)

[Compactness] $\boldsymbol{a} \mapsto \boldsymbol{g}(\boldsymbol{a})$ is continuous (i.e., $\mathcal{X}^*$ is compact)

[Convexity] for all players $i \in [n]$, and constraints $c \in [l]$ $\boldsymbol{a} \mapsto g_{ic}(\boldsymbol{a})$ is quasiconcave (i.e., $\mathcal{X}^*$ is convex)

Joint convexity of constraints is important as a VE is guaranteed to exist in all quasiconcave pseudo-games with jointly convex constraints. The proof of existence of VE was first provided by Rosen (1965). Analogous to the proof of existence of GNE, Rosen applies a fixed-point argument to a suitable best-response correspondence. We provide here a brief proof of existence and refer the reader to Facchinei and Kanzow (2010a) for additional context.

Theorem 9.2.2 [Theorem 1 of Rosen (1965)]**.**

A VE is guaranteed to exist in all quasiconcave pseudo-games with jointly convex constraints.

Proof of Theorem 9.2.1

Define the **VE best-response correspondence** as: $\mathcal{VBR}(\boldsymbol{a}) \doteq \arg\max_{\boldsymbol{a}'_i \in \mathcal{X}^*} \sum_{i \in [n]} u_i(\boldsymbol{a}'_i, \boldsymbol{a}_{-i})$. Now, consider any fixed point $\boldsymbol{a}^* \in \mathcal{VBR}(\boldsymbol{a}^*)$ of the VE best-response correspondence. Then, for all action profiles $\boldsymbol{a} \in \mathcal{X}^*$, and players $i \in [n]$, we have:

$$\sum_{i \in [n]} \left[u_i(\boldsymbol{a}_i, \boldsymbol{a}^*_{-i}) - u_i(\boldsymbol{a}^*) \right] \leq 0$$

That is, $\boldsymbol{a}^*$ is a VE. Now, by the maximum theorem (Berge, 1997), in quasiconcave games with jointly convex constraints, the VE best-response correspondence is guaranteed to be upper-hemicontinuous, non-empty-, compact-, and convex-valued. Hence, the best-response correspondence satisfies the assumptions of the Kakutani-

Glicksberg fixed point theorem (Theorem 2.4.1), and so a fixed point, and thus a VE, exists.

9.2.3 Nash Equilibrium and Generalized Nash Equilibrium Equivalence

While pseudo-games are a more practical modeling framework than games for mathematical analysis, for a large number of pseudo-games, the two models are equivalent. Excluding the rare cases where projection onto a constraint set can be computed in closed form, solution methods for constrained optimization problems often require solving an unconstrained penalized optimization problem. Borrowing this idea from optimization, for a large class of pseudo-games, it is possible to achieve a reduction from any n-player pseudo-game to a $2n$-player game, under the following standard constraint qualification.

Definition 9.2.6 [Slater's condition]**.**

A pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is said to satisfy **Slater's condition** iff for all $i \in [n]$, $c \in [l]$, and $\boldsymbol{a}_{-i} \in \mathcal{A}_{-i}$, there exists a **Slater vector** $\widetilde{\boldsymbol{a}}_i \in \operatorname{relint}(\mathcal{A}_i)$ s.t.:

$$\begin{aligned}
&\text{(Concave constraint function)} && \boldsymbol{a}_i \mapsto g_{ic}(\boldsymbol{a}_i, \boldsymbol{a}_{-i}) \text{ is concave} \\
&\text{(Weak Slater)} && \text{if } \boldsymbol{a}_i \mapsto g_{ic}(\boldsymbol{a}_i, \boldsymbol{a}_{-i}) \text{ is affine, then } g_{ic}(\tilde{\boldsymbol{a}}_i, \boldsymbol{a}_{-i}) \geq 0 \\
&\text{(Strong slater)} && \text{otherwise, } g_{ic}(\tilde{\boldsymbol{a}}_i, \boldsymbol{a}_{-i}) > 0
\end{aligned}$$

The following theorem shows that for any pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ that satisfies Slater's condition computing a GNE problem can be reduced to computing a NE, i.e., solving the following system of simultaneous penalized optimization problems:

$$\forall i \in [n], \quad \max_{\boldsymbol{a}_i \in \mathcal{A}_i} u_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i}) + \left\langle \boldsymbol{\lambda}_i, \boldsymbol{g}_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i}) \right\rangle \quad \forall i \in [n], \quad \min_{\boldsymbol{\lambda}_i \in \mathbb{R}^l_+} u_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i}) + \left\langle \boldsymbol{\lambda}_i, \boldsymbol{g}_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i}) \right\rangle$$

Note that in the above formulation, the objectives in the minimization problems should be interpreted as negated payoffs in the game.

Theorem 9.2.3 [Pseudo-game to game reduction]**.**

Consider a pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ that satisfies Slater's condition. Define the game $(2n, \mathcal{A}', \boldsymbol{u}')$, where for all $i \in [2n]$:

$$\text{(Action space)} \quad \mathcal{A}'_i \doteq \begin{cases} \mathcal{A}_i & \text{if } i \in [n] \\ \mathbb{R}^l_+ & \text{otherwise} \end{cases}$$

$$\text{(Payoffs)} \quad u'_i(\boldsymbol{a}, \boldsymbol{\lambda}) \doteq \begin{cases} u_i(\boldsymbol{a}) + \langle \boldsymbol{\lambda}_i, \boldsymbol{g}_i(\boldsymbol{a}) \rangle & \text{if } i \in [n] \\ -u_{i-n}(\boldsymbol{a}) - \left\langle \boldsymbol{\lambda}_i, \boldsymbol{g}_{i-n}(\boldsymbol{a}) \right\rangle & \text{otherwise} \end{cases}$$

Then, $(\boldsymbol{a}^*, \boldsymbol{\lambda}^*) \in \mathcal{A} \times \mathbb{R}^l_+$ is a NE of $(2n, \mathcal{A}', \boldsymbol{u}')$ iff $\boldsymbol{a}^*$ is a GNE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$.

Proof

($\Longrightarrow$): Let $(\boldsymbol{a}^*, \boldsymbol{\lambda}^*) \in \mathcal{A} \times \mathbb{R}^l_+$ be a NE of $(2n, \mathcal{A}', \boldsymbol{u}')$. Then, for all players $i \in [n]$, $(\boldsymbol{a}_i^*, \boldsymbol{\lambda}_i^*)$ is saddle point (i.e., NE) of the following min-max optimization problem:

$$\max_{\boldsymbol{a}_i \in \mathcal{A}_i} \min_{\boldsymbol{\lambda} \in \mathbb{R}^l_+} u_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i}^*) + \left\langle \boldsymbol{\lambda}_i, \boldsymbol{g}_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i}^*) \right\rangle$$

Hence, by the KKT theorem (Kuhn and Tucker, 1951), for all players $i \in [n]$, $\boldsymbol{a}_i^*$ is a solution of the optimization problem:

$$\max_{\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{a}_{-i})} u_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i}^*)$$

That is, $\boldsymbol{a}^*$ is a GNE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$.

($\Longleftarrow$): Let $\boldsymbol{a}^* \in \mathcal{X}(\boldsymbol{a}^*)$ be a GNE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$. That is, for all players $i \in [n]$, $\boldsymbol{a}^*$ is a solution of:

$$\max_{\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{a}_{-i})} u_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i}^*)$$

Since Slater's condition is satisfied, by the KKT theorem (Kuhn and Tucker, 1951), for all players $i \in [n]$, there exists $\boldsymbol{\lambda}_i^* \in \mathbb{R}^l_+$ s.t. $(\boldsymbol{a}_i^*, \boldsymbol{\lambda}_i^*)$ is a solution of the following Langrangian saddle-point problem:

$$\max_{\boldsymbol{a}_i \in \mathcal{A}_i} \min_{\boldsymbol{\lambda} \in \mathbb{R}^l_+} u_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i}^*) + \left\langle \boldsymbol{\lambda}_i, \boldsymbol{g}_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i}^*) \right\rangle$$

That is, $(\boldsymbol{a}^*, \boldsymbol{\lambda}^*) \in \mathcal{A} \times \mathbb{R}^l_+$ is a NE of $(2n, \mathcal{A}', \boldsymbol{u}')$.

Remark 9.2.3 [Boundedness of KKT multipliers]**.**

We note that under Slater's condition, any KKT multiplier $\boldsymbol{\lambda}^*$ associated with the NE

$(\boldsymbol{a}^*, \boldsymbol{\lambda}^*) \in \mathcal{A} \times \mathbb{R}_+^{nl}$ of the game $(2n, \mathcal{A}', \boldsymbol{u}')$ is necessarily bounded. In particular, by Lemma 3 of Nedic and Ozdaglar (2009), we can define a a non-empty, compact, convex set $\Lambda \subseteq \mathbb{R}_+^{nl}$, whose diameter depends on the value of the payoff functions of the players evaluated at the Slater vector of the game, which can be shown to contain the optimal KKT multipliers associated with the NE $(\boldsymbol{a}^*, \boldsymbol{\lambda}^*) \in \mathcal{A} \times \mathbb{R}_+^{nl}$ of the game $(2n, \mathcal{A}', \boldsymbol{u}')$.

With this reduction in hand, we will for the rest of this chapter focus on the computation of VE. Nevertheless, readers interested in applying the algorithms provided in this thesis can use the equivalence provided by the above theorem to compute a GNE in pseudo-games for which a VE does not exist.

9.3 Algorithms for Pseudo-Games

9.3.1 Computational model

With the question of existence answered, we now turn our attention to the computation of VE, and thus GNE (respectively, NE), in concave pseudo-games (respectively, games). While a GNE is guaranteed to exist in quasiconcave pseudo-games, we will restrict our attention to concave pseudo-games as the computation of an ε-NE even in single player quasiconcave games (i.e., quasiconcave optimization) is known to be NP-hard (Vavasis, 1995).

Algorithms for the computation of GNE, VE, and NE can be categorized into two main categories, decentralized algorithms, called **uncoupled learning dynamics**, in which each player employs a learning algorithm independently from the others, and centralized algorithms, which aim to compute an equilibrium without any restrictions (e.g., by a center, rather than by the players themselves).

Definition 9.3.1 [kth-order learning dynamics]**.**
Given some $k \in \mathbb{N}_{++}$, a pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ for which the derivatives $\{\nabla^j \boldsymbol{u}\}_{j=1}^{k-1}$ are well defined, and an initial iterate $\boldsymbol{a}^{(0)} \in \mathcal{A}$, a **$k$th-order learning dynamic $\boldsymbol{\pi}$** consists of an update function that generates the sequence of iterates $\{\boldsymbol{a}^{(t)}\}_t$ given by: for all $t = 0, 1, \ldots$,

$$\boldsymbol{a}^{(t+1)} \doteq \boldsymbol{\pi}\left(\bigcup_{i=0}^{t}(\boldsymbol{a}^{(i)}, \{\nabla^j \boldsymbol{u}(\boldsymbol{a}^{(i)})\}_{j=0}^{k-1})\right)$$

The most prominent class of kth-order learning dynamics in the literature on equilibrium computation, are special class of first-order learning dynamics for games called uncoupled learning dynamics (Hart and Mas-Colell, 2003) (for a recent survey, see for instance, Golowich et al. (2020a)).

Definition 9.3.2 [Uncoupled learning dynamics for games]**.**
Given a game $(\mathcal{A}, \boldsymbol{u})$, and an initial action profile $\boldsymbol{a} \in \mathcal{A}$, an **(first-order) uncoupled learning dynamic** (Hart and Mas-Colell, 2003) $\boldsymbol{\pi} = (\boldsymbol{\pi}_1, \ldots, \boldsymbol{\pi}_n)$ consists of an update function $\boldsymbol{\pi}_i : \bigcup_{\tau \geq 1}(\mathcal{A}_i \times \mathbb{R} \times \mathcal{A}_i^*) \to \mathcal{A}_i$ for each player $i \in [n]$, which generates the sequence

of actions $\{\boldsymbol{a}^{(t)}\}_t$ given by: for all $t = 0, 1, \ldots$,

$$\boldsymbol{a}_i^{(t+1)} \doteq \boldsymbol{\pi}_i \left(\bigcup_{k=0}^{t} (\boldsymbol{a}_i^{(t)}, u_i(\boldsymbol{a}^{(t)}), \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^{(t)})) \right)$$

Remark 9.3.1 [Uncoupled Learning Dynamics in Games and Pseudo-games]**.**
Uncoupled learning dynamics in games can be understood as players playing the game repeatedly, and updating their action at each round based on the observations in prior rounds without coordinating with other players, thus the "uncoupled" terminology. By the constrained nature of pseudo-games, it is in general inappropriate to consider uncoupled learning dynamics, since players are restricted to playing actions that are feasible w.r.t. one another's during repeated play of the pseudo-game.

As a remedy, one can generalize the notion of uncoupled learning dynamics for pseudo-games by introducing an "arbiter" who collects all players' action updates, and then projects them back onto the space of feasible action profiles $\mathcal{X}^*$. To avoid introducing heavy notation, we will not develop such machinery here; we simply note that the mirror extragradient learning dynamics we study for the computation of VE can be seen as such a type of an uncoupled learning dynamics. Further justifying this generalized definition, we note that when these first-order learning dynamics are instead applied to games, as we will show, the arising dynamics correspond to uncoupled learning dynamics. As a result, we will call a learning dynamic for pseudo-games "uncoupled" if, when applied to a game, the learning dynamics are uncoupled.

The computational complexity results in this chapter rely on the following computational model, which has been broadly adopted by the literature (see, for instance, Golowich et al. (2020a)).

Definition 9.3.3 [Pseudo-Game Computational Model]**.**
Given a pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ and a kth-order learning dynamic $\boldsymbol{\pi}$, the computational complexity of a kth-order learning dynamic is measured in term of the number of evaluations of the functions $\boldsymbol{u}, \nabla \boldsymbol{u}, \ldots, \nabla^k \boldsymbol{u}$.

Remark 9.3.2.

In line with the literature, the computational model we consider thus assumes that any other operation, such as (Bregman) projection onto a set, accrues a constant cost.

The computational results that exist in the literature, as well as the results we will present in this chapter, hold in the following two classes of pseudo-games.

Definition 9.3.4 [Lipschitz-Smooth Pseudo-Games]**.**
Given a modulus of smoothness $\lambda \geq 0$, a pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is said to be **λ-Lipschitz-smooth** iff for all players $i \in [n]$, $\nabla_{\boldsymbol{a}_i} u_i$ is λ-Lipschitz-continuous.

Definition 9.3.5 [Jointly Lipschitz-Smooth Pseudo-Games]**.**
Given a modulus of smoothness $\lambda \geq 0$, a **jointly λ-Lipschitz-smooth** pseudo-game is a pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ such that for all players $i \in [n]$, u_i is λ-Lipschitz-smooth.

9.3.2 Related Works

Following Arrow and Debreu's introduction of GNE, Rosen (1965) initiated the study of the mathematical and computational properties of GNE in pseudo-games with jointly convex constraints, proposing a projected gradient method to compute GNE. Thirty years later, Uryas'ev and Rubinstein (1994) developed the first relaxation methods for finding GNEs, which were improved upon in subsequent works (Krawczyk and Uryasev, 2000; Heusinger and Kanzow, 2009). Two other types of algorithms were also introduced to the literature: Newton-style methods (Facchinei et al., 2009; Dreves, 2017; Von Heusinger et al., 2012; Izmailov and Solodov, 2014; Fischer et al., 2016; Dreves et al., 2013) and interior-point potential methods (Dreves et al., 2013). Many of these approaches are based on minimizing the exploitability of the pseudo-game, but others use variational inequalities (Facchinei et al., 2007; Nabetani et al., 2011) and Lemke methods (Schiro et al., 2013).

More recently, novel methods that transform the problem of computing a GNE to that of a NE were analyzed. These models take the form of either exact penalization methods, which lift the constraints into the objective function via a penalty term (Facchinei and

Lampariello, 2011; Fukushima, 2011; Kanzow and Steck, 2018; Ba and Pang, 2020; Facchinei and Kanzow, 2010b), or augmented Lagrangian methods (Pang and Fukushima, 2005; Kanzow, 2016; Kanzow and Steck, 2018; Bueno et al., 2019), which do the same, augmented by dual Lagrangian variables. Using these methods, Jordan et al. (2022) provide the first convergence rates to an ε-GNE in monotone (respectively, strongly monotone) pseudo-games with jointly affine constraints in $\tilde{O}(1/\varepsilon)$ (respectively, $\tilde{O}(1/\sqrt{\varepsilon})$) iterations. These algorithms, despite being highly efficient in theory, are numerically unstable in practice (Jordan et al., 2022). Nearly all of the aforementioned approaches concerned pseudo-games with jointly convex constraints.

Exploitability minimization has also been a valuable tool in multi-agent reinforcement learning; algorithms in this literature that aim to minimize exploitability are known as exploitability-descent algorithms. Lockhart et al. (2019) analyzed exploitability descent in two-player, zero-sum, extensive-form games with finite action spaces. Variants of exploitability-descent have also been combined with entropic regularization and homotopy methods to solve for NE in large games (Gemp et al., 2021).

9.4 Computation of GNE

9.4.1 Uncoupled Learning Dynamics for GNE

The first type of algorithms we study for the computation of VE are uncoupled learning dynamics. For convenience, going forward, we will define the following important operator.

Definition 9.4.1 [Pseudo-Game operator]**.**
The **pseudo-game operator** associated with any pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is defined as

$$\boldsymbol{v}(\boldsymbol{a}) \doteq -(\nabla_{\boldsymbol{a}_1} u_1(\boldsymbol{a}), \ldots, \nabla_{\boldsymbol{a}_n} u_n(\boldsymbol{a}))$$

We denote the ith component of $\boldsymbol{v}(\boldsymbol{a})$ as $\boldsymbol{v}_i(\boldsymbol{a})$, which we note is equal to the negated gradient $-\nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a})$ of the ith player's payoff w.r.t. its own action.

Remark 9.4.1 [Generalization for subdifferentiable pseudo-games]**.**
In the above definition, for clarity it is assumed that $\nabla_{\boldsymbol{a}_i} u_i$ is well-defined. However, the definition can easily be extended to pseudo-games for which the subdifferential $\mathcal{D}_{\boldsymbol{a}_i} u_i$ is guaranteed to be non-empty, by defining the pseudo-game operator $\boldsymbol{v}$ as a correspondence, now better denoted $\mathcal{V}(\boldsymbol{a}) \doteq -\bigtimes_{i \in [n]} \mathcal{D}_{\boldsymbol{a}_i} u_i(\boldsymbol{a})$. The following lemma then also directly generalizes to such pseudo-games by replacing the VI $(\mathcal{X}^*, \boldsymbol{v})$ by the VI $(\mathcal{X}^*, \mathcal{V})$.

Be we introduce the uncoupled learning dynamic we will study, we present the following lemma, which uncovers a relationship between the VE of a pseudo-game and the strong solutions of the corresponding VI.

Lemma 9.4.1 [SVI $\subset$ VE in Concave Pseudo-Games]**.**
Given a concave pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$, any ε-strong solution of the VI $(\mathcal{X}^*, \boldsymbol{v})$ is an ε-VE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$.

Proof of Lemma 9.4.1

Let $\boldsymbol{a}^* \in \mathcal{X}^*$ be an ε-strong solution of the VI $(\mathcal{X}^*, \boldsymbol{v})$. Then, for all $\boldsymbol{a} \in \mathcal{X}^*$,

$$\begin{aligned}
\varepsilon &\geq \langle \boldsymbol{v}(\boldsymbol{a}^*), \boldsymbol{a}^* - \boldsymbol{a} \rangle \\
&= \sum_{i \in [n]} \left\langle -\nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i^* - \boldsymbol{a}_i \right\rangle \\
&= \sum_{i \in [n]} \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i - \boldsymbol{a}_i^* \right\rangle \\
&\geq \sum_{i \in [n]} u_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i}^*) - u_i(\boldsymbol{a}^*)
\end{aligned}$$

where the last line follows by concavity. Hence, $\boldsymbol{a}^*$ is an ε-VE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$.

In games, $\mathcal{X}^* = \mathcal{A}$, which implies that the set of VE is equal to the set of GNE, and hence the set of NE. Thus, we have the following corollary of Lemma 9.4.1.

Corollary 9.4.1 [SVI $\subseteq$ VE in Concave Games]**.**

Given a concave game $(\mathcal{A}, \boldsymbol{u})$, any ε-strong solution of the VI $(\mathcal{X}^*, \boldsymbol{v})$ is an ε-NE of $(\mathcal{A}, \boldsymbol{u})$.

With this lemma in hand, we can now apply the mirror extragradient method to solve the VI $(\mathcal{X}^*, \boldsymbol{v})$ associated with any concave pseudo-game $(\mathcal{A}, \boldsymbol{u})$ giving us the mirror extragradient learning dynamics (Algorithm 7). Here, we remark that when the pseudo-game considered is in fact a game, then the mirror extragradient method can be seen as an uncoupled learning dynamic.

Algorithm 7 Mirror Extragradient Learning Dyanmics

Input: $\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}, \tau, \eta, h, \boldsymbol{a}^{(0)}$

Output: $\{\boldsymbol{a}^{(t)}, \boldsymbol{a}^{(t+0.5)}\}_{t \in [\tau]}$

1: **for** $t = 1, \dots, \tau$ **do**

2: $\forall i \in [n], \quad \boldsymbol{a}_i^{(t+0.5)} \leftarrow \arg\max_{\boldsymbol{a}_i \in \mathcal{A}_i} \left\{ \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^{(t)}), \boldsymbol{a}_i - \boldsymbol{a}_i^{(t)} \right\rangle - \frac{1}{2\eta} \mathrm{div}_h(\boldsymbol{a}_i, \boldsymbol{a}_i^{(t)}) \right\}$

3: $\forall i \in [n], \boldsymbol{a}_i^{(t+1)} \leftarrow \arg\max_{\boldsymbol{a}_i \in \mathcal{A}_i} \left\{ \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^{(t+0.5)}), \boldsymbol{a}_i - \boldsymbol{a}_i^{(t)} \right\rangle - \frac{1}{2\eta} \mathrm{div}_h(\boldsymbol{a}_i, \boldsymbol{a}_i^{(t)}) \right\}$

return $\{\boldsymbol{a}^{(t)}, \boldsymbol{a}^{(t+0.5)}\}_{t \in [\tau]}$

Remark 9.4.2 [Mirror Extragradient Learning Dynamics Are Uncoupled]**.**
It turns out that the mirror extragradient learning dynamics (Algorithm 7), when applied to a game $(\mathcal{A}, \boldsymbol{u})$, can be seen as an uncoupled learning dynamic.

To see this, suppose that the mirror extragradient algorithm applied to $(\mathcal{A}, \boldsymbol{v})$ generates the sequence of action profiles $\{\boldsymbol{a}'^{(t)}, \boldsymbol{a}'^{(t+0.5)}\}_{t \in [0,\tau]}$. Now, for notational convenience, map the indices of the sequence through the transformation $t \mapsto 2t$ to obtain the sequence $\{\boldsymbol{a}^{(t)}\}_{t \in [0,2\tau]}$. The transformed sequence of actions can now be interpreted as an uncoupled learning dynamic where the update function for each player $i \in [n]$ and any $t \in \mathbb{N}_+$ is given by:

$$\boldsymbol{\pi}_i \left(\bigcup_{k=0}^{t} (\boldsymbol{a}_i^{(t)}, u_i(\boldsymbol{a}^{(t)}), \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^{(t)})) \right)$$
$$\doteq \begin{cases} \mathop{\arg\max}\limits_{\boldsymbol{a}_i \in \mathcal{A}_i} \left\{ \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^{(t)}), \boldsymbol{a}_i - \boldsymbol{a}_i^{(t)} \right\rangle - \frac{1}{2\eta} \mathrm{div}_h(\boldsymbol{a}_i, \boldsymbol{a}_i^{(t)}) \right\} & \text{if } t \text{ is even} \\ \mathop{\arg\max}\limits_{\boldsymbol{a}_i \in \mathcal{A}_i} \left\{ \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^{(t)}), \boldsymbol{a}_i - \boldsymbol{a}_i^{(t-1)} \right\rangle - \frac{1}{2\eta} \mathrm{div}_h(\boldsymbol{a}_i, \boldsymbol{a}_i^{(t-1)}) \right\} & \text{if } t \text{ is odd} \end{cases} .$$

We note that the $\arg\min$ is singleton-valued when the kernel function h is strictly convex, and hence we interpret the $\arg\min$ as the item in the singleton output. The arising update rule then dictates that on even time steps all players take a step of mirror ascent using their action the current time-step, while on odd time-steps, players take a step of mirror ascent using their action from the previous time-step.

With this in mind, we now introduce the class of variationally stable games for which we can prove the convergence of the mirror extragradient learning dynamics. While a definition of variational stability was first introduced by Zhou et al. (2017) for games, the definition we provide here is much weaker.[2] While Zhou et al. (2017) proved the asymptotic convergence of mirror ascent dynamics with decreasing step size in games when the kernel function h satisfies a set of regularity conditions, it is not clear if there exists a kernel function which satisfies these conditions and whether this convergence

[2]When generalized for pseudo-games, per Zhou et al., a pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is said to be variationally stable if its set of VE is equal to the set of weak solutions $\mathcal{MVI}(\mathcal{X}^*, \boldsymbol{v})$ of the VI $(\mathcal{X}^*, \boldsymbol{v})$. In contrast, to Zhou et al.'s definition, we only require the set of weak solutions of the VI $(\mathcal{X}^*, \boldsymbol{v})$ to be non-empty, hence generalizing Zhou et al.'s definition.

implies polynomial-time computation of ε-NE in such games (or more broadly of ε-VE in pseudo-games). In contrast, here, we provide a non-asymptotic convergence rate under very mild assumptions on the kernel function h, namely strong convexity, which implies the polynomial-time computation of ε-VE in variationally stable pseudo-games.

Definition 9.4.2 [Variationally Stable Pseudo-Games]**.**

A pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is said to be **variationally stable** iff there exists $\boldsymbol{a}^* \in \mathcal{X}^*$ s.t. for all $\boldsymbol{b} \in \mathcal{X}^*$:

$$\sum_{i \in [n]} \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{b}), \boldsymbol{a}_i^* - \boldsymbol{b}_i \right\rangle \geq 0$$

In other words, the pseudo-game is variationally stable iff the set of weak solutions $\mathcal{MVI}(\mathcal{X}^*, \boldsymbol{v})$ of the VI $(\mathcal{X}^*, \boldsymbol{v})$ is non-empty.

Remark 9.4.3 [Interpreting variational stability in concave pseudo-games]**.**

In concave games pseudo-games, by concavity, we have, for all $i \in [n]$, and $\boldsymbol{a}, \boldsymbol{b} \in \mathcal{A}$,

$$u_i(\boldsymbol{a}_i, \boldsymbol{b}_{-i}) - u_i(\boldsymbol{b}) \leq \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{b}), \boldsymbol{a}_i - \boldsymbol{b}_i \right\rangle$$

Hence, the variational stability condition is satisfied if there exists $\boldsymbol{a}^* \in \mathcal{X}^*$ s.t. for all $i \in [n]$ and $\boldsymbol{b} \in \mathcal{X}^*$,

$$u_i(\boldsymbol{a}_i^*, \boldsymbol{b}_{-i}) - u_i(\boldsymbol{b}) \geq 0$$

In other words, variational stability in concave pseudo-games can be ensured if for all players $i \in [n]$, there exists an action $\boldsymbol{a}_i^* \in \mathcal{X}^*$ that weakly increases its payoff when the player i unilaterally deviates from the action profile $\boldsymbol{b}$.

The class of variationally stable concave pseudo-games contains a number of well-studied pseudo-games games such as monotone pseudo-games with jointly convex constraints.[3]

Definition 9.4.3 [Monotone pseudo-games]**.**

A pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is said to be **monotone** iff the pseudo-game operator $\boldsymbol{v}$ is mono-

[3]Note that a pseudo-game being monotone implies that for all $i \in [n]$ and $\boldsymbol{a}_{-i} \in \mathcal{A}_{-i}$, $\boldsymbol{a}_i \mapsto u_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i})$ is concave. It does not, however, imply continuity of u_i, nor the non-emptiness, compactness, nor convexity of $\mathcal{X}^*$, all of which are necessary for the variational stability condition to hold.

tone, i.e., for all $\boldsymbol{a}, \boldsymbol{b} \in \mathcal{A}$,

$$\sum_{i \in [n]} \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}) - \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{b}), \boldsymbol{a}_i - \boldsymbol{b}_i \right\rangle \leq 0$$

The class of variationally stable concave pseudo-games also contains the larger but less well-studied class of pseudomonotone and quasimonotone concave pseudo-games with jointly convex constraints (see, for instance, Section 2.3.2 of Huang and Zhang (2023)):

Definition 9.4.4 [Pseudomonotone games]**.**

A pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is said to be **pseudomonotone** iff the pseudo-game operator $\boldsymbol{v}$ is pseudomonotone, i.e., for all $\boldsymbol{a}, \boldsymbol{b} \in \mathcal{A}$,

$$\sum_{i \in [n]} \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}), \boldsymbol{b}_i - \boldsymbol{a}_i \right\rangle \leq 0 \implies \sum_{i \in [n]} \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{b}), \boldsymbol{b}_i - \boldsymbol{a}_i \right\rangle \leq 0$$

Definition 9.4.5 [Quasimonotone Pseudo-Games]**.**

A pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is said to be **quasimonotone** iff the pseudo-game operator is quasimonotone $\boldsymbol{v}$, i.e., for all $\boldsymbol{a}, \boldsymbol{b} \in \mathcal{A}$,

$$\sum_{i \in [n]} \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}), \boldsymbol{b}_i - \boldsymbol{a}_i \right\rangle < 0 \implies \sum_{i \in [n]} \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{b}), \boldsymbol{b}_i - \boldsymbol{a}_i \right\rangle \leq 0$$

With these definitions in hand, we can obtain a non-asymptotic convergence rate for the mirror extragradient learning dynamics in variationally stable pseudo-games, as an application of Theorem 4.3.1 in conjunction with Lemma 9.4.1.

Theorem 9.4.1 [Convergence of Mirror Extragradient Learning Dynamics]**.**

Let $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ be a variationally stable and λ-Lipschitz-smooth concave pseudo-game with jointly convex constraints and h a 1-strongly-convex and κ-Lipschitz-smooth kernel function. Consider the mirror extragradient learning dynamics (Algorithm 7) run with the pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$, the kernel function h, a step size $\eta \in \left(0, \frac{1}{\sqrt{2}\lambda}\right]$, and any time horizon $\tau \in \mathbb{N}$. The output sequence $\{\boldsymbol{x}^{(t+0.5)}, \boldsymbol{x}^{(t+1)}\}_t$ satisfies: If $\boldsymbol{a}_{\text{best}}^{(\tau)} \in \arg\min_{\boldsymbol{x}^{(k+0.5)}:k=0,\dots,\tau} \operatorname{div}_h(\boldsymbol{a}^{(k+0.5)}, \boldsymbol{a}^{(k)})$, then for some $\tau \in O(1/\varepsilon^2)$, $\boldsymbol{a}_{\text{best}}^{(\tau)}$ is an ε-VE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$. In addition, the iterates asymptotically converge to a VE $\boldsymbol{a}^* \in \mathcal{X}^*$ of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$, i.e., $\lim_{t\to\infty} \boldsymbol{a}^{(t+0.5)} = \lim_{t\to\infty} \boldsymbol{a}^{(t)} = \boldsymbol{a}^*$.

Proof of Theorem 9.4.1

By Lemma 9.4.1, we know that any ε-strong solution of the VI $(\mathcal{X}^*, \boldsymbol{v})$ is a ε-VE of the pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$. Now, note that by the variational stability assumption, the set of weak solutions of $(\mathcal{X}^*, \boldsymbol{v})$ is non-empty. In addition, as $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is a jointly λ-Lipschitz-smooth concave pseudo-game with jointly convex constraints, $(\mathcal{X}^*, \boldsymbol{v})$ is λ-Lipschitz continuous. Hence, the assumptions of Theorem 4.3.1 hold, giving us the result.

With this theorem in hand, so remarks are in order.

Remark 9.4.4 [Contributions to the literature]**.**

To the best of our knowledge, Theorem 9.4.1 is the broadest polynomial-time computation result for ε-VE in pseudo-games, as well as ε-NE in games. It is also the first and only existing non-asymptotic convergence analysis of the mirror extragradient learning dynamics.

We note that for the choice of kernel function $h \doteq \|\cdot\|^2$ (i.e., the Euclidean squared norm), while Huang and Zhang (2023) do not explicitly prove the above result, it could be inferred from their Theorem 3.16, when taken in conjuction with Lemma 9.4.1. As such, in this specific setting, our contribution can be seen as identifying Lemma 9.4.1 as applicable to Huang and Zhang's result for pseudo-games.

Remark 9.4.5 [Local convergence to ε-VE]**.**

The above finite-time global convergence result to ε-VE, can be extended to a finite-time local convergence result to ε-VE by instead applying Theorem 4.3.2 with the assumption that the initial iterate of the mirror extragradient learning dynamics starts close enough to a local weak solution of $(\mathcal{X}^*, \boldsymbol{v})$. To the best of our knowledge this is the first finite-time local convergence result to ε-VE in pseudo-games, as well as the first finite-time local convergence result to ε-NE in games.

9.4.2 Merit Function Methods for GNE

We now turn our attention to methods with computational guarantees beyond variationally stable pseudo-games. In general concave pseudo-games, it is not possible for uncoupled learning dynamics to converge a VE (see, Theorem 1 of Hart and Mas-Colell (2003)). To remedy this non-convergence issue, we will in this section consider consider first-order *coupled* learning dynamics, or simply first-order learning dynamics. To derive, these first-order learning dynamics, we will define merit functions for VE and consider methods to minimize this merit functions.

Definition 9.4.6 [Merit functions for Pseudo-Games]**.**

Given a pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$. A function $\Xi : \mathcal{X}^* \to \mathbb{R}$ is said to be a **merit function** for the set of VE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ iff

1. for all $\boldsymbol{a} \in \mathcal{X}^*$, $\Xi(\boldsymbol{a}) \geq 0$
2. for any $\boldsymbol{a}^* \in \mathcal{X}^*$, $\Xi(\boldsymbol{a}^*) = 0$ iff $\boldsymbol{a}^*$ is a VE.

Our formulations start with the exploitability, or the Nikaido-Isoda function (Nikaido and Isoda, 1955), as well as the related cumulative regret or Ky Fan function (Aubin, 2013) of a pseudo-game.

Definition 9.4.7 [Cumulative Regret and Exploitability]**.**

The **cumulative regret** (or **Ky Fan function**) $\psi : \mathcal{A} \times \mathcal{A} \to \mathbb{R}$ between two action profiles $\boldsymbol{a}$ and $\boldsymbol{b}$ across all players is defined as:

$$\psi(\boldsymbol{a}, \boldsymbol{b}) \doteq \sum_{i \in [n]} \left[u_i(\boldsymbol{b}_i, \boldsymbol{a}_{-i}) - u_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i}) \right]$$

The **exploitability**, or the **Nikaido-Isoda *potential* function** (Nikaido and Isoda, 1955), $\varphi : \mathcal{A} \to \mathbb{R}$ of an action profile $\boldsymbol{a}$ is defined as

$$\varphi(\boldsymbol{a}) = \max_{\boldsymbol{b} \in \mathcal{X}^*} \sum_{i \in [n]} \left[u_i(\boldsymbol{b}_i, \boldsymbol{a}_{-i}) - u_i(\boldsymbol{a}_i, \boldsymbol{a}_{-i}) \right] \quad .$$

It is well known that any unexploitable action profile, i.e., $\boldsymbol{a} \in \mathcal{X}(\boldsymbol{a})$ s.t. $\varphi(\boldsymbol{a}) = 0$, in a pseudo-game is a VE.

Lemma 9.4.2 [(Flam and Ruszczynski, 1994)].

Given a pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$, it holds that for all $\boldsymbol{a} \in \mathcal{X}^*$, $\varphi(\boldsymbol{a}) \geq 0$. Moreover, any action profile $\boldsymbol{a} \in \mathcal{X}(\boldsymbol{a})$ with exploitability $\varphi(\boldsymbol{a})$ is an $\varphi(\boldsymbol{a})$-VE. Further, any action profile $\boldsymbol{a}^{(*)} \in \mathcal{X}^*$ is a VE iff it achieves the lowerbound, i.e., $\varphi(\boldsymbol{a}^*) = 0$.

This lemma tells us that we can reformulate the VE computation problem as the optimization problem of minimizing exploitability, i.e., $\min_{\boldsymbol{a} \in \mathcal{X}^*} \varphi(\boldsymbol{a})$. Despite this reformulation of VE computation in terms of exploitability, no convergence rate guarantees are known for exploitability-minimization algorithms. This unexploitability (!) of exploitability may be due to the fact that it is not differentiable in general. The key insight that allows us to obtain convergence guarantees is that we treat the VE problem not as a minimization problem, but rather as a min-max optimization problem, namely:

$$\min_{\boldsymbol{a} \in \mathcal{X}^*} \varphi(\boldsymbol{a}) = \min_{\boldsymbol{a} \in \mathcal{X}^*} \max_{\boldsymbol{b} \in \mathcal{X}^*} \psi(\boldsymbol{a}, \boldsymbol{b})$$

This problem is well understood when ψ is a convex-concave objective function (Nemirovski, 2004; Korpelevich, 1976; Nedic and Ozdaglar, 2009; von Neumann, 1928). Furthermore, the cumulative regret ψ is indeed convex-concave, i.e., convex in $\boldsymbol{a}$ and concave in $\boldsymbol{b}$, in many pseudo-games of interest: e.g., two-player zero-sum, n-player pairwise zero-sum, and a large class of monotone and bilinear pseudo-games, as well as Cournot oligopoly games, to name a few. For more details, see Section 2 of Flam and Ruszczynski (1994).

Using the simple observation that every VE of a pseudo-game is the solution to a min-max optimization problem we introduce our first algorithm (EDA; Algorithm 8), an extragradient method (Korpelevich, 1976). The algorithm works by interleaving extragradient ascent and descent steps: at iteration t, given $\boldsymbol{a}^{(t)}$, it ascends on $\psi(\boldsymbol{a}^{(t)}, \cdot)$, thereby generating a better response $\boldsymbol{b}^{(t+1)}$, and then it descends on $\psi(\cdot, \boldsymbol{b}^{(t+1)})$, thereby decreasing exploitability. We combine several known results about the convergence of extragradient descent methods in min-max optimization problems to obtain the following convergence guarantees for EDA in pseudo-games.

Algorithm 8 Extragradient descent ascent (EDA)

Inputs: $\psi, \tau, \eta, \boldsymbol{a}^{(0)}, \boldsymbol{b}^{(0)}$

Outputs: $(\boldsymbol{a}^{(t)}, \boldsymbol{b}^{(t)}, \boldsymbol{a}^{(t+0.5)}, \boldsymbol{b}^{(t+0.5)})_t$

1: **for** $t = 0, \dots, \tau - 1$ **do**

2: $\quad \boldsymbol{a}^{(t+0.5)} = \Pi_{\mathcal{X}^*}\left[\boldsymbol{a}^{(t)} - \eta \nabla_{\boldsymbol{a}} \psi(\boldsymbol{a}^{(t)}, \boldsymbol{b}^{(t)})\right]$

3: $\quad \boldsymbol{b}^{(t+0.5)} = \Pi_{\mathcal{X}^*}\left[\boldsymbol{b}^{(t)} + \eta \nabla_{\boldsymbol{b}} \psi(\boldsymbol{a}^{(t)}, \boldsymbol{b}^{(t)})\right]$

4: $\quad \boldsymbol{a}^{(t+1)} = \Pi_{\mathcal{X}^*}\left[\boldsymbol{a}^{(t)} - \eta \nabla_{\boldsymbol{a}} \psi(\boldsymbol{a}^{(t+0.5)}, \boldsymbol{b}^{(t+0.5)})\right]$

5: $\quad \boldsymbol{b}^{(t+1)} = \Pi_{\mathcal{X}^*}\left[\boldsymbol{b}^{(t)} + \eta \nabla_{\boldsymbol{b}} \psi(\boldsymbol{a}^{(t+0.5)}, \boldsymbol{b}^{(t+0.5)})\right]$

6: **return** $(\boldsymbol{a}^{(t)}, \boldsymbol{b}^{(t)}, \boldsymbol{a}^{(t+0.5)}, \boldsymbol{b}^{(t+0.5)})_t$

Remark 9.4.6 [From EDA to Mirror Extragradient Descent Ascent]**.**

We note that one could more generally consider a mirror extragradient descent ascent method, which is equivalent to running the mirror extragradient algorithm (Algorithm 3) on the VI $(\mathcal{X}^* \times \mathcal{X}^*, (\nabla_{\boldsymbol{a}} \psi, -\nabla_{\boldsymbol{b}} \psi))$. The convergence results we provide in this section hold for this more general algorithm by simply applying Theorem 4.3.1 and other related theorems in the literature (e.g., Nemirovski (2004)). For simplicity, we choose to present our results for this simpler algorithm.

Theorem 9.4.2 [Convergence of EDA]**.**

Consider a jointly λ-Lipschitz-smooth quasiconcave pseudo-game with jointly convex-constraints $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ with convex-concave cumulative regret ψ. Suppose that EDA (Algorithm 8) is run with the cumulative regret ψ, the step size $\eta \leq \frac{1}{2n\lambda}$, time horizon $\tau \in \mathbb{N}$, and initial iterates $\boldsymbol{a}^{(0)}, \boldsymbol{b}^{(0)} \in \mathcal{X}^*$. The output sequence $(\boldsymbol{a}^{(t)}, \boldsymbol{b}^{(t)}, \boldsymbol{a}^{(t+0.5)}, \boldsymbol{b}^{(t+0.5)})_t$ satisfies the following: If $\overline{\boldsymbol{a}}^{(\tau)} = {}^{1}\!/_{\tau} \sum_{t=1}^{T} \boldsymbol{a}^{(t)}$ and $\boldsymbol{a}_{\text{best}}^{(\tau')} \in \underset{\boldsymbol{a}^{(k+0.5)}: k=0,1,\dots,\tau'}{\arg\min} \|\boldsymbol{a}^{(k+0.5)} - \boldsymbol{a}^{(k)}\|$, then for all $\varepsilon \geq 0$, there exists $\tau \in O\left({}^{n\lambda}\!/_{\varepsilon}\right)$ and $\tau' \in O\left({}^{n\lambda}\!/_{\varepsilon^2}\right)$ s.t. $\overline{\boldsymbol{a}}^{(\tau)}$ and $\boldsymbol{a}_{\text{best}}^{(\tau')}$ are both ε-VE.

Proof of Theorem 9.4.2

First, note that ψ is $2n\lambda$-Lipschitz-smooth since the pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is jointly λ-Lipschitz-smooth, and ψ is the sum of n differences of λ-Lipschitz-smooth functions. Additionally, since ψ is convex-concave, $(\mathcal{X}^* \times \mathcal{X}^*, (\nabla_{\boldsymbol{a}} \psi, -\nabla_{\boldsymbol{b}} \psi))$ is monotone. Further, note that any feasible action profile $\boldsymbol{a}^* \in \mathcal{X}^*$ is an ε-strong solution of $(\mathcal{X}^* \times \mathcal{X}^*, (\nabla_{\boldsymbol{a}} \psi, -\nabla_{\boldsymbol{b}} \psi))$ iff $\max_{\boldsymbol{b} \in \mathcal{X}^*} \psi(\boldsymbol{a}^*, \boldsymbol{b}) - \min_{\boldsymbol{a} \in \mathcal{X}^*} \max_{\boldsymbol{b} \in \mathcal{X}^*} \psi(\boldsymbol{a}, \boldsymbol{b}) \leq \varepsilon$ (see, Proposition 2.2 of Nemirovski (2004)). Since a VE is guaranteed to exist under the assumption of joint convexity (Theorem 9.2.2), it holds that $\min_{\boldsymbol{a} \in \mathcal{X}^*} \max_{\boldsymbol{b} \in \mathcal{X}^*} \psi(\boldsymbol{a}, \boldsymbol{b}) = 0$. Therefore,

$$
\begin{aligned}
\varepsilon &\geq \max_{\boldsymbol{b} \in \mathcal{X}^*} \psi(\boldsymbol{a}^*, \boldsymbol{b}) - \underbrace{\min_{\boldsymbol{a} \in \mathcal{X}^*} \max_{\boldsymbol{b} \in \mathcal{X}^*} \psi(\boldsymbol{a}, \boldsymbol{b})}_{=0} \\
&= \max_{\boldsymbol{b} \in \mathcal{X}^*} \psi(\boldsymbol{a}^*, \boldsymbol{b}) \\
&= \varphi(\boldsymbol{a}^*)
\end{aligned}
$$

As such, by Lemma 9.4.2, any ε-strong solution of $(\mathcal{X}^* \times \mathcal{X}^*, (\nabla_{\boldsymbol{a}} \psi, -\nabla_{\boldsymbol{b}} \psi))$ is an ε-VE.

Hence, by Theorem 3.2 of Nemirovski (2004), for all $\varepsilon \geq 0$ and $\tau \geq {}^{2n\lambda}/_{\varepsilon}$, $\overline{\boldsymbol{a}}^{(\tau)}$ is an ε-VE.

Similarly, by Theorem 4.3.1, for all $\varepsilon \geq 0$, there exists $\tau' \in O({}^{n\lambda}/_{\varepsilon})$ s.t. $\boldsymbol{a}_{\text{best}}^{(\tau')}$ is an ε-VE.

We note that EDA is an optimal algorithm for computing VE in pseudo-games with convex-concave cumulative regret.

Remark 9.4.7 [Optimality of EDA]**.**

The computational complexity of two-player zero-sum convex-concave games is $\Omega({}^{1}/_{\varepsilon})$. Since pseudo-games with convex-concave cumulative regret are a special case, the iteration complexity of EDA in pseudo-games with convex-concave cumulative regret is optimal.

We also remark that this result can be extended beyond pseudo-games in which the cumulative regret is convex-concave. We present this less general result for the sake of simplicity, and because it answers an open question first posed by Flam and Ruszczynski (1994), who asked whether exploitability can be minimized efficiently in pseudo-games when ψ is convex-concave.

Remark 9.4.8 [EDA under Minty Condition]**.**
Going beyond pseudo-games in which ψ is convex-concave, we can consider pseudo-games in which the VI $(\mathcal{X}^* \times \mathcal{X}^*, (\nabla_{\boldsymbol{a}} \psi, -\nabla_{\boldsymbol{b}} \psi))$ satisfies the Minty condition. For such pseudo-games, by Theorem 4.3.1, it is still possible compute an ε-VE in polynomial time: for all $\varepsilon \geq 0$, there exists $\tau' \in O(n\lambda/\varepsilon)$ s.t. $\boldsymbol{a}_{\text{best}}^{(\tau')}$ is an ε-VE.

Remark 9.4.9 [EDA in concave pseudo-games with jointly convex constraints]**.**
In general, for concave pseudo-games with jointly convex constraints, while the cumulative regret ψ is not necessarily convex-convex, it is guaranteed to be non-convex-concave. In this more general setting, as exploitability is not differentiable, it is not possible to show convergence to an ε-stationary point of φ. Nevertheless, it is possible to show that extragradient descent ascent can compute an ε-stationary point of the Moreau envelope $\tilde{\varphi}$ of φ, defined as $\tilde{\varphi}(\boldsymbol{a}) \doteq \min_{\boldsymbol{b} \in \mathcal{X}^*} \varphi(\boldsymbol{b}) + 1/2n\lambda \|\boldsymbol{a} - \boldsymbol{b}\|^2$ in $O(1/\varepsilon^6)$ operations (see, for instance Mahdavinia et al. (2022)).

While Remark 9.4.9 suggests that it at least possible to minimize exploitability in general concave pseudo-games with jointly convex constraints, as ε-stationary points of $\tilde{\varphi}$ are not directly related to ε-stationary points of φ, it seems hard to relate this convergence result to the computation of ε-VE. Nonetheless, by exploiting the structure of cumulative regret, we regularize it to obtain a regularized exploitability function whose set of minima is once again equal to the set of VE.

In particular, observe the following: if $\boldsymbol{a}^* \in \arg\min_{\boldsymbol{a} \in \mathcal{X}} \max_{\boldsymbol{b} \in \mathcal{X}} \psi(\boldsymbol{a}, \boldsymbol{b})$, then $\boldsymbol{a}^* \in \arg\max_{\boldsymbol{b} \in \mathcal{X}} \psi(\boldsymbol{a}^*, \boldsymbol{b})$. In other words, if $\boldsymbol{a}^*$ is a solution to the outer minimization problem, then it is likewise a solution to the inner maximization problem. As a result, we can penalize

Algorithm 9 Regularized Extragradient Descent Ascent (REDA)

Inputs: $\psi_\alpha, \tau, \eta, \boldsymbol{a}^{(0)}, \boldsymbol{b}^{(0)}$

Outputs: $(\boldsymbol{a}^{(t)}, \boldsymbol{b}^{(t)}, \boldsymbol{a}^{(t+0.5)}, \boldsymbol{b}^{(t+0.5)})_t$

1: **for** $t = 0, \ldots, \tau - 1$ **do**

2: $\quad \boldsymbol{a}^{(t+0.5)} = \Pi_{\mathcal{X}^*}\left[\boldsymbol{a}^{(t)} - \eta \nabla_{\boldsymbol{a}} \psi_\alpha(\boldsymbol{a}^{(t)}, \boldsymbol{b}^{(t)})\right]$

3: $\quad \boldsymbol{b}^{(t+0.5)} = \Pi_{\mathcal{X}^*}\left[\boldsymbol{b}^{(t)} + \eta \nabla_{\boldsymbol{b}} \psi_\alpha(\boldsymbol{a}^{(t)}, \boldsymbol{b}^{(t)})\right]$

4: $\quad \boldsymbol{a}^{(t+1)} = \Pi_{\mathcal{X}^*}\left[\boldsymbol{a}^{(t)} - \eta \nabla_{\boldsymbol{a}} \psi_\alpha(\boldsymbol{a}^{(t+0.5)}, \boldsymbol{b}^{(t+0.5)})\right]$

5: $\quad \boldsymbol{b}^{(t+1)} = \Pi_{\mathcal{X}^*}\left[\boldsymbol{b}^{(t)} + \eta \nabla_{\boldsymbol{b}} \psi_\alpha(\boldsymbol{a}^{(t+0.5)}, \boldsymbol{b}^{(t+0.5)})\right]$

6: **return** $(\boldsymbol{a}^{(t)}, \boldsymbol{b}^{(t)}, \boldsymbol{a}^{(t+0.5)}, \boldsymbol{b}^{(t+0.5)})_t$

exploitability in proportion to the distance between $\boldsymbol{a}$ and $\boldsymbol{b}$, while still ensuring that this penalized exploitability is minimized at a VE.

We thus propose to optimize the α-**regularized cumulative regret** $\psi_\alpha : \mathcal{A} \times \mathcal{A} \to \mathbb{R}$,

Definition 9.4.8 [Regularized Cumulative Regret and Regularized Exploitability]**.**
Let $\alpha \geq 0$ be a regularization parameter.

The α-**regularized cumulative regret** of any two action profiles $\boldsymbol{a}, \boldsymbol{b} \in \mathcal{A}$ is defined as:

$$\psi_\alpha(\boldsymbol{a}, \boldsymbol{b}) \doteq \psi(\boldsymbol{a}, \boldsymbol{b}) - \frac{\alpha}{2} \left\| \boldsymbol{a} - \boldsymbol{b} \right\|_2^2$$

The α-**regularized exploitability** $\varphi_\alpha : \mathcal{A} \to \mathbb{R}$ of any action is defined as:

$$\varphi_\alpha(\boldsymbol{a}) \doteq \max_{\boldsymbol{b} \in \mathcal{X}} \psi_\alpha(\boldsymbol{a}, \boldsymbol{b})$$

Von Heusinger and Kanzow show that an action profile $\boldsymbol{a}^*$ has no α-regularized-exploitability, i.e., $\varphi_\alpha(\boldsymbol{a}^*) = 0$, iff $\boldsymbol{a}^*$ is a VE, for all $\alpha \geq 0$ (Theorem 3.3 of Von Heusinger and Kanzow (2009)). With this observation in hand, we can then try to minimize the regularized exploitability by running extragradient descent ascent on ψ_α, rather than ψ, which gives us the regularized extragradient descent ascent algorithm (Algorithm 9).

Theorem 9.4.3 [Convergence of REDA]**.**

Consider a jointly λ-Lipschitz-smooth concave pseudo-game with jointly convex-constraints $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$. Let $\alpha > 0$ be some regularization parameter, and ψ_α be the α-regularized cumulative regret associated with $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$. Suppose that REDA (Algorithm 9) is run with is run with the regularized cumulative regret ψ_α, the step size $\eta \leq \min\{\frac{1}{75\alpha(2n\lambda+\alpha)}, \frac{1}{4(2n\lambda+\alpha)}\}$, time horizon $\tau \in \mathbb{N}$, and initial iterates $\boldsymbol{a}^{(0)}, \boldsymbol{b}^{(0)} \in \mathcal{X}^*$. The output sequence $(\boldsymbol{a}^{(t)}, \boldsymbol{b}^{(t)}, \boldsymbol{a}^{(t+0.5)}, \boldsymbol{b}^{(t+0.5)})_t$ satisfies the following: If $\boldsymbol{a}^{(\tau)}_{\text{best}} \in \arg\min_{\boldsymbol{a}^{(k)}:k=0,\dots,\tau-1} \max_{\boldsymbol{a}\in\mathcal{X}^*} \langle \nabla\varphi_\alpha(\boldsymbol{a}^{(k)}), \boldsymbol{a}^{(k)} - \boldsymbol{a}\rangle$, then for all $\varepsilon \geq 0$, there exists $\tau \in O(1/\varepsilon^2)$, $\boldsymbol{a}^{(\tau)}_{\text{best}}$ is an ε-stationary point of φ_α, i.e., $\max_{\boldsymbol{a}\in\mathcal{X}^*} \langle \nabla\varphi_\alpha(\boldsymbol{a}^{(\tau)}_{\text{best}}), \boldsymbol{a}^{(\tau)}_{\text{best}} - \boldsymbol{a}\rangle \leq \varepsilon$.

Proof of Theorem 9.4.3

The result follows from an application of Theorem 4.2 Mahdavinia et al. (2022) by noting the following. First, note that Theorem 4.2 is derived for min-max optimization problems where the minimization is unconstrained, which in our application corresponds to the optimization problem $\min_{\boldsymbol{a}\in\mathbb{R}^{n\times m}} \max_{\boldsymbol{b}\in\mathcal{X}^*} \psi(\boldsymbol{a}, \boldsymbol{b})$. Nevertheless, as Mahdavinia et al. (2022) remark at the end of their related work section, their proof directly generalizes to the constrained setting, with a definition of an ε-stationary point of φ_α given by $\max_{\boldsymbol{x}\in\mathcal{X}} \langle \nabla\varphi_\alpha(\boldsymbol{a}^{(\tau)}_{\text{best}}), \boldsymbol{a}^{(\tau)}_{\text{best}} - \boldsymbol{a}\rangle \leq \varepsilon$.
Second, for all $\boldsymbol{a} \in \mathcal{A}$, $\boldsymbol{b} \mapsto \psi_\alpha(\boldsymbol{a}, \boldsymbol{b})$ is α is α-strongly-concave, since it is the sum of a concave function and an α-strongly-concave function.
Finally, note that ψ_α is $(2n\lambda + \alpha)$-Lipschitz-smooth, since it the sum of n differences of λ functions ψ, and $\frac{\alpha}{2}\|\boldsymbol{a} - \boldsymbol{b}\|_2^2$ which is α-Lipschitz-smooth. Hence, setting the step size so that $\eta \leq \min\{\frac{1}{75\alpha(2n\lambda+\alpha)}, \frac{1}{4(2n\lambda+\alpha)}\}$, the antecedant of Theorem 4.2 of Mahdavinia et al. (2022) is satisfied, and the result follows.

With this convergence result in hand, one might wonder under what conditions ε-stationary points of ψ_α coincide with ε-VE.

Remark 9.4.10 [When is an ε-stationary point of φ_α also an ε-VE]**.**

Von Heusinger and Kanzow (2009) show in Theorem 3.6 that ε-stationary points of ψ_α

coincide with ε-VE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$[4], when the pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ satisfies the following strict monotonicity-like condition: for all $\boldsymbol{a}, \boldsymbol{b}$ s.t. $\boldsymbol{a} \neq \boldsymbol{b}$,

$$\sum_{i \in [n]} \left\langle \nabla u_i(\boldsymbol{b}_i, \boldsymbol{a}_{-i}) - \nabla u_i(\boldsymbol{a}), \boldsymbol{a} - \boldsymbol{b} \right\rangle > 0$$

This condition can be translated to a more intuitive condition on $\boldsymbol{u}$ by noting that it is satisfied when for all players $i \in [n]$, $\boldsymbol{a} \mapsto u_i(\boldsymbol{a})$ is strictly convex.

Remark 9.4.11 [Extending Convergence to Quasiconcave Pseudo-Games]**.**
We note that while the cumulative regret function for all $\boldsymbol{a} \in \mathcal{A}$, $\boldsymbol{b} \mapsto \psi(\boldsymbol{a}, \boldsymbol{b})$ is not concave in quasiconcave pseudo-games, under the assumption that the pseudo-game is jointly λ-Lipschitz-smooth, $\boldsymbol{b} \mapsto \psi(\boldsymbol{a}, \boldsymbol{b})$ is λ-weakly-concave. Hence, choosing $\alpha > \lambda$, we can ensure that $\boldsymbol{b} \mapsto \psi_\alpha(\boldsymbol{a}, \boldsymbol{b})$ is strongly-concave, hence allowing us to extend Theorem 9.4.3 to jointly λ-Lipschitz-smooth quasiconcave pseudo-games with jointly convex constraints.

Unfortunately, beyond quasiconcave pseudo-games with jointly convex constraints, a VE is not guaranteed to exist, and as such minimizing the (regularized) exploitability is no longer a sensible approach. Nevertheless, it is possible to show the existence of a weaker solution concept in smooth pseudo-games, namely first-order variational equilibrium. As such, we next turn our attention to the computation of first-order variational equilibrium.

[4]Von Heusinger and Kanzow (2009) show the result for 0-stationary points and 0-VE; however, their proof directly generalizes to ε-stationary points by replacing the 0s in their proof with εs.

9.5 Local Solution Concepts and Existence

9.5.1 First-Order and Local Generalized Nash and Variational Equilibrium

We now turn our attention to solving non-concave games. As previously mentioned, the computation of even an ε-NE in single player quasiconcave games (i.e., quasiconcave optimization) is known to be NP-hard (Vavasis, 1995). As such, to obtain any computational results for efficient algorithms, we focus on local solutions.

Further, as many pseudo-games that arise in modern machine learning applications are not quasiconcave, and as GNE are not guaranteed to exist beyond quasiconcave games, this begs the question: how far beyond quasiconcave games can the theory of games be extended, and what solution concepts are appropriate for such games (and, more generally, pseudo-games)? Two intuitive solution concepts which are weaker than GNE are the first-order generalized Nash equilibrium and the local generalized Nash equilibrium.

Definition 9.5.1 [First-Order GNE]**.**

Given $\varepsilon \geq 0$, an ε-**first-order generalized Nash equilibrium** (ε-**first-order GNE**) is an action profile $\boldsymbol{a}^* \in \mathcal{X}(\boldsymbol{a}^*)$ s.t. for all players $i \in [n]$ and $\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{a}^*_{-i})$,

$$\left\langle \partial_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i - \boldsymbol{a}^*_i \right\rangle \leq \varepsilon \ \ ,$$

for some $\partial_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*) \in \mathcal{D}_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*)$.

A 0-first-order GNE is simply called a first-order generalized Nash equilibrium (first-order GNE).

Remark 9.5.1 [Interpretation of first-order GNE]**.**

First-order GNE can be interpreted as the GNE of the pseudo-game with "linearized" payoffs around the GNE. More precisely, for all players $i \in [n]$, let $\ell_i^{\boldsymbol{a}^*}[u_i](\boldsymbol{a}_i) \doteq u_i(\boldsymbol{a}^*) + \left\langle \partial_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i - \boldsymbol{a}^*_i \right\rangle$ be the linearization operator around $\boldsymbol{a}^*$ s.t. $\ell_i^{\boldsymbol{a}^*}[u_i](\boldsymbol{a}_i)$ is a first-order Taylor expansion approximation of $u_i(\boldsymbol{a}_i, \boldsymbol{a}^*_{-i})$.

Recall that a GNE is an action profile $\boldsymbol{a}^* \in \mathcal{X}(\boldsymbol{a}^*)$ s.t. for all $i \in [n]$ and $\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{a}^*_{-i})$,

$$u_i(\boldsymbol{a}^*) \geq u_i(\boldsymbol{a}_i, \boldsymbol{a}^*_{-i}) \ \ .$$

Now, suppose we replace the payoffs $\boldsymbol{u}$ in the original pseudo-game with the linearized payoffs around $\boldsymbol{a}^*$. Then the above definition reduces to: for all $i \in [n]$ and $\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{a}^*_{-i})$,

$$\begin{aligned}
&\ell_i^{\boldsymbol{a}^*}[u_i](\boldsymbol{a}_i^*) \geq \ell_i^{\boldsymbol{a}^*}[u_i](\boldsymbol{a}_i) \\
\iff & u_i(\boldsymbol{a}^*) \geq u_i(\boldsymbol{a}^*) + \left\langle \partial_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i - \boldsymbol{a}_i^* \right\rangle \\
\iff & 0 \geq \left\langle \partial_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i - \boldsymbol{a}_i^* \right\rangle
\end{aligned}$$

This interpretation is key in the analysis of many general equilibrium models in modern macroeconomics, because in most if not all models analyzed in practice, the cumulative expected utility of the consumers is linearized before solving the model (Sargent and Ljungqvist, 2000; Auclert et al., 2021). As such, the first-order GNE provides a theoretical framework to understand this trick that is so prevalently used in practice.

We can similarly define a first-order analog of variational equilibrium, which is a refinement of the first-order GNE.

Definition 9.5.2 [First-Order VE]**.**

Given $\varepsilon \geq 0$, an ε-**first-order variational equilibrium** (ε-**first-order VE**) is an action profile $\boldsymbol{a}^* \in \mathcal{X}(\boldsymbol{a}^*)$ s.t. for all $\boldsymbol{a} \in \mathcal{X}^*$,

$$\sum_{i \in [n]} \left\langle \partial_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i - \boldsymbol{a}_i^* \right\rangle \leq \varepsilon \ ,$$

for some $\partial_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*) \in \mathcal{D}_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*)$.

A 0-first-order VE is simply called a first-order variational equilibrium (first-order VE).

Note that the set of first-order VE of any pseudo-game is a subset of the set of first-order GNE. An alternative to first-order GNE is local GNE.

Definition 9.5.3 [Local GNE]**.**

Given a **regret parameter** $\varepsilon \geq 0$, and a **locality parameter** $\delta \geq 0$, a (ε, δ)**-local generalized Nash equilibrium** ((ε, δ)**-local GNE**) is an action profile $\boldsymbol{a}^* \in \mathcal{A}$ s.t. for all players $i \in [n]$ and $\boldsymbol{a}_i \in \mathcal{A}_i \cap \mathcal{B}_\delta[\boldsymbol{a}_i^*]$,

$$u_i(\boldsymbol{a}^*) \geq u_i(\boldsymbol{a}_i, \boldsymbol{a}^*_{-i}) - \varepsilon \ .$$

For any $\delta > 0$, a $(0, \delta)$-local GNE is simply called a local δ-generalized Nash equilibrium (local δ-GNE).

Similarly, we can define a local analog of VE.

Definition 9.5.4 [Local VE]**.**
Given a **regret parameter** $\varepsilon \geq 0$, and a **locality parameter** $\delta \geq 0$, a (ε, δ)**-local variational equilibrium** ((ε, δ)**-local VE**) is an action profile $\boldsymbol{a}^* \in \mathcal{A}$ s.t. for all $\boldsymbol{a} \in \mathcal{X}^* \cap \mathcal{B}_\delta[\boldsymbol{a}^*]$,

$$\sum_{i \in [n]} u_i(\boldsymbol{a}_i, \boldsymbol{a}^*_{-i}) - u_i(\boldsymbol{a}^*) \leq \varepsilon \ .$$

For any $\delta > 0$, a $(0, \delta)$-local VE is called a local δ-variational equilibrium (local δ-VE).

Similar to their global variants, the set of local VE is a subset of the set of local GNE.

Remark 9.5.2 [First-Order NE and Local NE in Games]**.**
In games, first-order GNE and first-order VE are equivalent, so these solution concepts are simply called **first-order Nash equilibrium** (**first-order NE**). Similarly, the definitions of local GNE and local VE are also equivalent, so these solution concepts are simply called **local Nash equilibrium** (**local NE**).

9.5.2 Smooth Pseudo-Games

Unfortunately, the existence of local GNE in even very simply non-concave games cannot be guaranteed.

Example 9.5.1 [Exact local GNE non-existence]**.**
Consider the two player zero-sum game $(2, \mathcal{A}, \boldsymbol{u})$ where $\mathcal{A} \doteq [-1, 1] \times [-1, 1]$, and $u_1(a_1, a_2) = -u_2(a_1, a_2) = (a_1 - a_2)^2$. At a local Nash equilibrium $(a_1^*, a_2^*) \in \mathcal{A}$, the max player would play $a_1 \neq a_2$, while the min player would play $a_1 = a_2$. Hence, an exact local NE cannot exist.

Nevertheless for any $\varepsilon, \delta \geq 0$ s.t. $\varepsilon \geq 2\delta^2$, any action profile $(a_1^*, a_2^*) \in \mathcal{A}$ s.t. $a_1^* = a_2^*$ is an (ε, δ)-local NE since:

$$\begin{aligned}
&\max_{a_1 \in [-1,1] \cap \mathcal{B}_\delta[a_1^*]} u_1(a_1, a_2^*) && \max_{a_2 \in [-1,1] \cap \mathcal{B}_\delta[a_2^*]} u_2(a_1^*, a_2) \\
&= \max_{a_1 \in [-1,1] \cap \mathcal{B}_\delta[a_2^*]} (a_1 - a_2^*)^2 && \max_{a_2 \in [-1,1] \cap \mathcal{B}_\delta[a_2^*]} -(a_2^* - a_2)^2 \\
&\leq \delta^2 && = 0
\end{aligned}$$

In contrast, first-order GNE can be shown to exist in smooth pseudo-games.

Definition 9.5.5.

A **smooth pseudo-game** is a pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$, where for all players $i \in [n]$:

[Continuous payoffs]	$\nabla_{\boldsymbol{a}_i} u_i$ is continuous
[Convex constraints]	$\mathcal{X}_{-i}$ is continuous, non-empty-, compact-, and convex-valued
[Convex action space]	$\mathcal{A}_i$ is non-empty, compact, and convex

Theorem 9.5.1 [Existence of first-order GNE]**.**

A first-order GNE exists in any smooth pseudo-game.

Proof of Theorem 9.5.1

Consider the first-order best-response correspondence $\mathcal{FOBR}(\boldsymbol{a}) \doteq \bigtimes_{i \in [n]} \arg\min_{\boldsymbol{a}_i' \in \mathcal{X}_i(\boldsymbol{a}_{-i})} \left\{ \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}), \boldsymbol{a}_i' \right\rangle \right\}$. Note that at any fixed point $\boldsymbol{a}^*$ s.t. $\boldsymbol{a}^* \in \mathcal{FOBR}(\boldsymbol{a}^*)$, for all players $i \in [n]$,

$$\boldsymbol{a}_i^* \in \arg\min_{\boldsymbol{a}_i' \in \mathcal{X}_i(\boldsymbol{a}_{-i}^*)} \left\{ \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i' \right\rangle \right\}$$

Equivalently, for all $i \in [n]$ and $\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{a}_{-i}^*)$,

$$\left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i^* \right\rangle \leq \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i \right\rangle \iff \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i^* - \boldsymbol{a}_i \right\rangle \leq 0$$

Hence, $\boldsymbol{a}^*$ is first-order GNE.

Now, note that by the Berge's maximum theorem Berge (1997), in smooth games, the first-order best-response correspondence $\mathcal{FOBR}$ is upper hemicontinuous, and non-empty-, compact-, convex-valued. Hence, by the Kakutani-Glicksberg fixed point theorem (Theorem 2.4.1) a fixed point of $\mathcal{FOBR}$ exists, and so a first-order GNE is likewise guaranteed to exist.

A similar existence result can be shown for VE, under the additional assumption that the pseudo-game has jointly convex constraints.

Theorem 9.5.2 [Existence of first-order VE]**.**
A first-order VE exists in any smooth pseudo-game with jointly convex constraints.

Proof of Theorem 9.5.2

Consider the variational first-order best-response correspondence $\mathcal{VFOBR}(\boldsymbol{a}) \doteq \arg\min_{\boldsymbol{a}' \in \mathcal{X}^*} \left\{ \sum_{i \in [n]} \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}), \boldsymbol{a}'_i \right\rangle \right\}$. Note that at any fixed point $\boldsymbol{a}^*$ s.t. $\boldsymbol{a}^* \in \mathcal{VFOBR}(\boldsymbol{a}^*)$, for all players $i \in [n]$,

$$\boldsymbol{a}^* \in \arg\min_{\boldsymbol{a}'_i \in \mathcal{X}^*} \left\{ \sum_{i \in [n]} \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}'_i \right\rangle \right\}$$

Equivalently, for all players $i \in [n]$ and $\boldsymbol{a} \in \mathcal{X}^*$,

$$\sum_{i \in [n]} \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i^* \right\rangle \leq \sum_{i \in [n]} \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i \right\rangle \iff \sum_{i \in [n]} \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i^* - \boldsymbol{a}_i \right\rangle \leq 0$$

Hence, $\boldsymbol{a}^*$ is first-order VE.

Now, note that by the Berge's maximum theorem Berge (1997), in smooth games with jointly convex constraints, the variational first-order best-response correspondence $\mathcal{VFOBR}$ is upper hemicontinuous, and non-empty-, compact-, convex-valued. Hence, by the Kakutani-Glicksberg fixed point theorem (Theorem 2.4.1) a fixed point of $\mathcal{VFOBR}$ exists, and so a first-order VE is likewise guaranteed to exist.

9.5.3 First-Order and Local Equilibrium Equivalence

While a local GNE is not guaranteed to exist in general, in λ-Lipschitz-smooth pseudo-games, we can show that for any $\varepsilon \geq 0$, an ε-first-order GNE is an $(\varepsilon + {}^{\lambda\delta^2}/_2, \delta)$-local GNE, hence guaranteeing the existence of a $(\lambda\delta^2/2, \delta)$-local GNE, for any choice of $\delta \geq 0$, by the existence of first-order GNE.

Lemma 9.5.1 [First-Order and Local GNE Equivalence]**.**
Given a λ-Lipschitz-smooth game, there exists $\varepsilon, \delta \geq 0$ s.t. any ε-first-order GNE (respectively, VE) is an $(\varepsilon + \lambda\delta^2/2, \delta)$-local GNE (respectively, $(\varepsilon + \lambda\delta^2/2, \delta)$-local VE).

Proof of Lemma 9.5.1

Assume an ε-first-order GNE $\boldsymbol{a}^* \in \mathcal{A}$. By the weak-concavity property of Lipschitz-smooth functions, for all $i \in [n]$ and $\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{a}^*_{-i})$,

$$\begin{aligned}
u_i(\boldsymbol{a}_i, \boldsymbol{a}^*_{-i}) &\leq u_i(\boldsymbol{a}^*) + \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i - \boldsymbol{a}^*_i \right\rangle + \frac{\lambda}{2} \|\boldsymbol{a}_i - \boldsymbol{a}^*_i\|^2 \\
u_i(\boldsymbol{a}_i, \boldsymbol{a}^*_{-i}) &\leq u_i(\boldsymbol{a}^*) + \varepsilon + \frac{\lambda}{2} \|\boldsymbol{a}_i - \boldsymbol{a}^*_i\|^2 \\
u_i(\boldsymbol{a}_i, \boldsymbol{a}^*_{-i}) - u_i(\boldsymbol{a}^*) &\leq \varepsilon + \frac{\lambda}{2} \|\boldsymbol{a}_i - \boldsymbol{a}^*_i\|^2 \quad ,
\end{aligned}$$

where the penultimate line follows from the definition of an ε-first-order GNE.

Now, because $\boldsymbol{a}_i$ lies within a δ-ball around $\boldsymbol{a}^*_i$, it follows that for all $i \in [n]$ and $\boldsymbol{a} \in \mathcal{X}_i(\boldsymbol{a}_{-i}) \cap \mathcal{B}_\delta[\boldsymbol{a}^*_i]$,

$$u_i(\boldsymbol{a}_i, \boldsymbol{a}^*_{-i}) - u_i(\boldsymbol{a}^*) \leq \varepsilon + \frac{\lambda\delta^2}{2} \tag{9.1}$$

That is, $\boldsymbol{a}^*$ is an $(\varepsilon + \lambda\delta^2/2, \delta)$-local Nash equilibrium.

The proof follows similarly for the case of first-order VE and local VE.

9.6 Computation of First-Order and Local GNE

9.6.1 First-Order Variational Equilibrium and Variational Inequalities

We first present a characterization of first-order VE in pseudo-games in terms of strong solutions of VIs. The following intuitive generalization of Lemma 9.4.1 is guaranteed to hold for ε-first-order VE in all pseudo-games. We present this result for smooth pseudo-games only, as the computational results we will derive will require the payoff functions of the players to be differentiable. Nevertheless, the result directly generalizes to all pseudo-games by replacing the VI in the statement of Lemma 9.6.1 with the one presented in Remark 9.4.1.

Lemma 9.6.1 [SVI = VE in Smooth Pseudo-Games]**.**
The set of ε-first-order VE of any smooth pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is equal to the set of ε-strong solutions $\mathcal{SVI}_\varepsilon(\mathcal{X}^*, \boldsymbol{v})$ of the VI $(\mathcal{X}^*, \boldsymbol{v})$.

Proof of Lemma 9.6.1

Let $\boldsymbol{a}^* \in \mathcal{X}^*$ be a first-order VE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$. Then, for all $i \in [n]$ and $\boldsymbol{a} \in \mathcal{X}^*$, we have:

$$
\begin{aligned}
& \varepsilon \geq \sum_{i \in [n]} \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}^*), \boldsymbol{a}_i - \boldsymbol{a}_i^* \right\rangle \\
\iff \; & \varepsilon \geq \langle -\boldsymbol{v}(\boldsymbol{a}^*), \boldsymbol{a} - \boldsymbol{a}^* \rangle \\
\iff \; & \varepsilon \geq \langle \boldsymbol{v}(\boldsymbol{a}^*), \boldsymbol{a}^* - \boldsymbol{a} \rangle
\end{aligned}
$$

Hence, $\boldsymbol{a}^*$ is an ε-strong solution of the VI $(\mathcal{X}^*, \boldsymbol{v})$. Since the inequalities are equivalent, the converse is also true.

Since the set of first-order VE is equal to the set of first-order NE in games, and $\mathcal{X}^* = \mathcal{A}$, we also have the following corollary of Lemma 9.6.1.

Corollary 9.6.1.

The set of ε-NE of any smooth game $(\mathcal{A}, \boldsymbol{u})$ is equal to the set of ε-strong solutions of the VI $(\mathcal{A}, \boldsymbol{v})$.

9.6.2 Uncoupled Learning Dynamics First-Order GNE

Having characterized first-order VE in terms of the strong solutions of a VI, we now turn our attention to the computation of first-order VE. Beyond variationally stable concave pseudo-games, first-order methods are not guaranteed to converge to VE nor to GNE.[5] Thus, a very natural question to ask is what convergence guarantees can be obtained in variationally stable pseudo-games that are not necessarily concave.

The class of variationally stable pseudo-games contains the class of quasimonotone (and hence, pseudomonotone and monotone) smooth pseudo-games with jointly convex constraints (see, for instance Huang and Zhang (2023)). Unfortunately, the class of quasimonotone games does not take us beyond the class of quasiconcave games as the following remark describes.

Remark 9.6.1 [Quasimonotone Pseudo-Games are Quasiconcave]**.**
We note that any quasimonotone game is quasiconcave. To see this, first recall that a function is quasiconvex iff its subdifferential is quasimonotone (see, for instance, Theorem 4.1 of (Aussel et al., 1994)).

Now, for all $k \neq i \in [n]$, setting $\boldsymbol{b}_k \doteq \boldsymbol{a}_k$, the quasimonotonicity condition implies: for all $i \in [n]$ and $\boldsymbol{a}_i, \boldsymbol{b}_i \in \mathcal{A}_i$,

$$\left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{a}), \boldsymbol{a}_i - \boldsymbol{b}_i \right\rangle < 0 \implies \left\langle \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{b}_i, \boldsymbol{a}_{-i}), \boldsymbol{a}_i - \boldsymbol{b}_i \right\rangle \leq 0$$

Hence for all $i \in [n]$ and $\boldsymbol{a}_{-i} \in \mathcal{A}_{-i}$, the mapping $\boldsymbol{b}_i \mapsto \nabla_{\boldsymbol{a}_i} u_i(\boldsymbol{b}_i, \boldsymbol{a}_{-i})$ is quasimonotone, implying that $\boldsymbol{b}_i \mapsto u_i(\boldsymbol{b}_i, \boldsymbol{a}_{-i})$ is quasiconcave.

[5] See, for instance Example 9.6.1, and observe that first-order methods can converge to $(0, 0)$, which is not a VE.

Nevertheless, while the class of quasimonotone pseudo-games contains the class quasiconcave pseudo-games, variationally stable games are not necessarily quasiconcave, as shown by the following example.

Example 9.6.1 [Variationally stable game that is not quasiconcave]**.**

Consider the single player game $(1, \mathcal{A}, \boldsymbol{u})$, where $\mathcal{A}_1 = [-1, 1]^2$ and $u_1 \doteq (\boldsymbol{a}_1) \doteq {}^1\!/_3\, a_{11}^3 + {}^1\!/_3\, a_{12}^3$, equivalently stated as the following constrained optimization problem:

$$\max_{a_{11}, a_{12} \in [-1,1]} u_i(\boldsymbol{a}_1) \doteq \frac{1}{3} a_{11}^3 + \frac{1}{3} a_{12}^3$$

For this game, we have $\nabla u_1(\boldsymbol{a}_1) \doteq (a_{11}^2, a_{12}^2)$. Now, notice that this game is variationally stable, since for $\boldsymbol{a}_1^* \doteq (1, 1)$, we have ,for all $\boldsymbol{b}_1 \in [-1, 1]$, $\nabla u_1(\boldsymbol{b}_1) \geq 0$, and hence $\langle \nabla u_1(\boldsymbol{b}), \boldsymbol{a}_i^* - \boldsymbol{b}_i \rangle \geq 0$. However, notice that u_1 is not quasiconcave, since $u_1({}^1\!/_2(1, 0) + {}^1\!/_2({}^1\!/_2, 1)) = u_1({}^3\!/_4, {}^1\!/_2) = \frac{35}{192} \approx 0.182$, but $u_1(1, 0) = \frac{1}{3} \approx 0.333$ and $u_1({}^1\!/_2, 1) = \frac{3}{8} \approx 0.375$, meaning that $u_1({}^1\!/_2(1, 0) + {}^1\!/_2({}^1\!/_2, 1)) < \min\{u_1(1, 0), u_1({}^1\!/_2, 1)\}$.

This observation suggests that the class of variational stable pseudo-games is an interesting and broad enough class of non-quasiconcave pseudo-games to be worthy of our study. Further, by Lemma 9.6.1, since variational stability ensures that the corresponding VI satisfies the Minty condition, it is also a sufficient condition to ensure the convergence of first-order methods. In particular, recall that by applying the mirror extragradient method to solve the corresponding VI, we obtained the mirror extragradient learning dynamics (Algorithm 7). Hence, applying Theorem 4.3.1 in conjuction with Lemma 9.6.1, we obtain the following convergence theorem for the mirror extragradient learning dynamics in variationally stable pseudo-games with jointly convex constraints.

Theorem 9.6.1 [Convergence of Mirror extragradient Learning Dynamics]**.**

Let $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ be a variationally stable and λ-Lipschitz-smooth pseudo-game with jointly convex constraints, and h a 1-strongly-convex and κ-Lipschitz-smooth kernel function. Consider the mirror extragradient learning dynamics (Algorithm 7) run on the pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ with the kernel function h, a step size $\eta \in \left(0, \frac{1}{\sqrt{2}\lambda}\right]$, for any time horizon $\tau \in \mathbb{N}$. The output sequence $\{\boldsymbol{x}^{(t+0.5)}, \boldsymbol{x}^{(t+1)}\}_t$ satisfies the following: If

$\boldsymbol{a}_{\text{best}}^{(\tau)} \in \arg\min_{\boldsymbol{x}^{(k+0.5)}:k=0,\ldots,\tau} \operatorname{div}_h(\boldsymbol{a}^{(k+0.5)}, \boldsymbol{a}^{(k)})$, then for some $\tau \in O(1/\varepsilon^2)$, $\boldsymbol{a}_{\text{best}}^{(\tau)}$ is an ε-first-order VE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$. In addition, the iterates asymptotically converge to a first-order VE $\boldsymbol{a}^* \in \mathcal{X}^*$ of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$, i.e., $\lim_{t\to\infty} \boldsymbol{a}^{(t+0.5)} = \lim_{t\to\infty} \boldsymbol{a}^{(t)} = \boldsymbol{a}^*$

Proof of Theorem 9.6.1

By Lemma 9.6.1, any ε-strong solution of the VI $(\mathcal{X}^*, \boldsymbol{v})$ is an ε-first-order VE of the smooth pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$. Now, note that by the variational stability assumption, the set of weak solutions of $(\mathcal{X}^*, \boldsymbol{v})$ is non-empty. In addition, as $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is a jointly λ-Lipschitz-smooth concave pseudo-games with jointly convex constraints, $(\mathcal{X}^*, \boldsymbol{v})$ is λ-Lipschitz continuous. Hence, the assumptions of Theorem 4.3.1 hold, giving us the result.

With this theorem in hand, we make two remarks on its implications on the local VE (or GNE), and the local convergence properties of the algorithm.

Remark 9.6.2 [Computation of Local GNE/VE]**.**
Applying Lemma 9.5.1 under the assumptions of Theorem 9.6.1, we can show that for all $\varepsilon, \delta \geq 0$ s.t. $\varepsilon \geq \lambda\delta^2/2$, there exists some choice of $\tau \in O(\frac{1}{\varepsilon^2})$, $\boldsymbol{a}_{\text{best}}^{(\tau)}$ is a (ε, δ)-local VE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$. Choices of (ε, δ) s.t. $\varepsilon \geq \lambda\delta^2/2$ have previously been known under the name of local parameter regimes (see, for instance (Daskalakis et al., 2020b) and (Daskalakis, 2022)), and do not contradict the non-existence of a local-VE, since under this choice of parameters $\varepsilon = 0$ iff $\delta = 0$.

Remark 9.6.3 [Local Convergence to ε-VE]**.**
As mentioned in Remark 9.4.5, the above finite-time global convergence result to ε-first-order VE can be extended to a finite-time local convergence result to ε-first-order VE by instead applying Theorem 4.3.2, under the assumption that the initial iterate of the mirror extragradient learning dynamics starts close enough to a local weak solution of the VI $(\mathcal{X}^*, \boldsymbol{v})$. To the best of our knowledge this is the first finite-time local convergence result to ε-first-order VE in pseudo-games, as well as ε-first-order NE in games.

Remark 9.6.4 [Contributions to the literature]**.**

To the best of our knowledge, the above result is the first polynomial-time computation result for ε-first-order VE, as well as ε-first-order NE, in variationally stable pseudo-games and games, respectively. It is also the first and only existing non-asymptotic convergence analysis of the mirror extragradient learning dynamics in such games.

9.6.3 Merit Function Methods for First-Order GNE

Unfortunately, as the following example shows, beyond variationally stable pseudo-games, it is not in general possible to guarantee the convergence of the mirror extragradient method to first-order GNE.

Example 9.6.2 [Non-Convergence of First-Order Methods Beyond Variationally Stable Pseudo-Games]**.**

Consider the two player zero-sum game $(2, \mathcal{A}, \boldsymbol{u})$ where $\mathcal{A} \doteq \mathbb{R} \times \mathbb{R}$ and $u_1(a_1, a_2) = -u_2(a_1, a_2) = (a_2 - a_1)^2$. The set of first-order VE of this game are given by $\{(\boldsymbol{a}_1, \boldsymbol{a}_2) \in \mathbb{R} \times \mathbb{R} \mid \boldsymbol{a}_1 = \boldsymbol{a}_2\}$. However, for any $a_1^{(0)} > a_2^{(0)}$, for any choice of step sizes, the iterates generated by the mirror extragradient learning dynamics tend to negative infinity, while for $a_1^{(0)} < a_2^{(0)}$, the iterates tend to infinity.

To overcome this non-convergence issue, we will instead consider second-order methods. Our approach to derive a second-order method method for pseudo-games will be to optimize a merit function associated with the first-order VE of the pseudo-game.[6] In particular, recall that by Lemma 9.6.1, the set of first-order VE of the pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ can be expressed as the set of strong solutions of the VI $(\mathcal{X}^*, \boldsymbol{v})$. As such, we will consider optimizing the regularized primal gap function associated with $(\mathcal{X}^*, \boldsymbol{v})$, which we call the variational exploitability.

Definition 9.6.1 [Variational Exploitability]**.**

Given a regularization parameter $\alpha \geq 0$, the α-variational exploitability $\Xi_\alpha : \mathcal{X}^* \to \mathbb{R}$ of any pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ is defined as:

$$\Xi_\alpha(\boldsymbol{a}) \doteq \max_{\boldsymbol{b} \in \mathcal{X}^*} \langle \boldsymbol{v}(\boldsymbol{a}), \boldsymbol{a} - \boldsymbol{b} \rangle - \frac{\alpha}{2} \|\boldsymbol{b} - \boldsymbol{a}\|^2 \tag{9.2}$$

The 0-variational exploitability is simply called the variational exploitability.

We note the following corollary of Lemma 4.4.1 for the variational exploitability, which confirms that it is a merit function for first-order VE.

[6] The definition of merit functions for VE intuitively extends to first-order VE, by replacing any mentions of "VE" in Definition 9.4.6 with "first-order-VE".

Corollary 9.6.2 [Properties of the regularized primal gap]**.**

Consider a continuous VI $(\mathcal{X}^*, \boldsymbol{v})$. If $\alpha > 0$, then $\max_{\boldsymbol{b} \in \mathcal{X}^*} \langle \boldsymbol{v}(\boldsymbol{a}), \boldsymbol{a} - \boldsymbol{b} \rangle - \frac{\alpha}{2} \|\boldsymbol{b} - \boldsymbol{a}\|^2$ has a unique solution. In addition, the following holds:

1. $\boldsymbol{b}^*(\boldsymbol{a}) = \arg\max_{\boldsymbol{b} \in \mathcal{X}^*} \langle \boldsymbol{v}(\boldsymbol{a}), \boldsymbol{a} - \boldsymbol{b} \rangle - \frac{\alpha}{2} \|\boldsymbol{b} - \boldsymbol{a}\|^2 \doteq \Pi_{\mathcal{X}^*} \left[\boldsymbol{a} - \frac{1}{\alpha} \boldsymbol{v}(\boldsymbol{a})\right]$
2. $\nabla \Xi_\alpha(\boldsymbol{a}) = \boldsymbol{v}(\boldsymbol{a}) - (\nabla \boldsymbol{v}(\boldsymbol{a}) + \alpha \mathbb{I}) (\boldsymbol{b}^*(\boldsymbol{a}) - \boldsymbol{a})$
3. $\Xi_\alpha(\boldsymbol{a}) = \max_{\boldsymbol{b} \in \mathcal{X}^*} \frac{\alpha}{2} \left[\left\|\frac{1}{\alpha} \boldsymbol{v}(\boldsymbol{a})\right\|^2 - \left\|\boldsymbol{b} - \left(\boldsymbol{a} - \frac{1}{\alpha} \boldsymbol{v}(\boldsymbol{a})\right)\right\|^2\right]$
4. For all $\boldsymbol{a} \in \mathcal{X}^*$, $\Xi_\alpha(\boldsymbol{a}) \geq 0$ and $\Xi_\alpha(\boldsymbol{a}) = 0$ iff $\boldsymbol{a}$ is first-order VE

Now, notice that we can minimize variational exploitability via a mirror descent method, but as the gradient $\nabla \Xi_\alpha$ involves $\boldsymbol{v}(\boldsymbol{a})$ and $\nabla \boldsymbol{v}(\boldsymbol{a})$, each of which respectively depends on the gradient and the Hessian of the players' utility functions, the corresponding method, which we call the mirror variational learning dynamics (Algorithm 10), is a second-order learning dynamic.

Algorithm 10 Mirror Variational Learning Dynamics

Input: $\Xi_\alpha, h, \tau, \eta, \boldsymbol{a}^{(0)}$

Output: $\{\boldsymbol{a}^{(t)}\}_t$

1: **for** $t = 1, \ldots, \tau$ **do**

2: $\quad \boldsymbol{a}^{(t+1)} \leftarrow \arg\min_{\boldsymbol{a} \in \mathcal{X}^*} \left\{\left\langle \nabla \Xi_\alpha(\boldsymbol{a}^{(t)}), \boldsymbol{a} - \boldsymbol{a}^{(t)} \right\rangle + \frac{1}{2\eta} \mathrm{div}_h(\boldsymbol{a}, \boldsymbol{a}^{(t)}) \right\}$

return $\{\boldsymbol{a}^{(t)}\}_t$

Now, since the mirror variational learning dynamic algorithm is an instance of the mirror potential method (Algorithm 5) applied to the VI $(\mathcal{X}^*, \boldsymbol{v})$, we obtain the following theorem as an application of Theorem 4.4.1.

Theorem 9.6.2 [Mirror Variational Learning Dynamics Convergence]**.**

Let $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ be a jointly convex, λ-Lipschitz-smooth pseudo-game with $\nabla^2 u_i$ β-Lipschitz-continuous for all $i \in [n]$, h a 1-strongly-convex kernel function, $\alpha \geq 0$, $\eta \in \left(0, \frac{1}{2(2\beta\alpha \mathrm{diam}(\mathcal{X}^*)^2 + 1 + 2\lambda)}\right]$, and $\boldsymbol{a}^{(0)} \in \mathcal{X}^*$. Consider the mirror variational learning dynamics (Algorithm 10) run on $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$ with the variational exploitability Ξ_α associated with $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u})$, the kernel function h, an arbitrary time horizon $\tau \in \mathbb{N}$, the step size η, and the initial iterate $\boldsymbol{a}^{(0)}$. The output sequence $\{\boldsymbol{a}^{(t)}\}_t$ converges to a stationary point of Ξ_α, respecting the following bound:

$$\min_{k=0,1,\ldots,\tau-1} \max_{\boldsymbol{a} \in \mathcal{X}^*} \langle \nabla \Xi_\alpha(\boldsymbol{a}^{(k)}), \boldsymbol{a}^{(k)} - \boldsymbol{a} \rangle \leq \frac{2\Xi_\alpha(\boldsymbol{a}^{(0)})}{\tau}$$

In addition, if $\boldsymbol{a}_{\text{best}}^{(\tau)} \in \arg\min_{\boldsymbol{a}^{(k)}:k=0,\ldots,\tau-1} \max_{\boldsymbol{a} \in \mathcal{X}^*} \langle \nabla \Xi_\alpha(\boldsymbol{a}^{(k)}), \boldsymbol{a}^{(k)} - \boldsymbol{a} \rangle$, then, for some $\tau \in O(1/\varepsilon)$, $\boldsymbol{a}_{\text{best}}^{(\tau)}$ is a ε-stationary point of Ξ_α.

With this theorem in hand, we conclude with a remark on its interpretation, before turning to some applications of our results.

Remark 9.6.5 [When are stationary points global solutions]**.**
As a corollary of Remark 4.4.1, we note that the stationary points of the variational exploitability correspond to VE when the pseudo-game is monotone. As such, Theorem 9.6.2 implies that a VE can be computed in polynomial time via the mirror variational learning dynamics in monotone pseudo-games.

Chapter 10

Arrow-Debreu Economies

10.1 Background

An **Arrow-Debreu economy** $(n, m, \mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, denoted $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ when n and m are clear from context, comprises a finite set of $m \in \mathbb{N}_+$ divisible **commodities** and $n \in \mathbb{N}_+$ **consumers**. Each consumer $i \in [n]$ is characterized by a **set of consumptions** $\mathcal{X}_i \subseteq \mathbb{R}^m$, an **endowment** of commodities $\boldsymbol{e}_i = (e_{i1}, \ldots, e_{im}) \in \mathbb{R}^n$, and a **utility function** $u_i : \mathbb{R}^m \to \mathbb{R}$, which describes the utility $u_i(\boldsymbol{x}_i)$ consumer i derives from **consumption** $\boldsymbol{x}_i \in \mathcal{X}_i$.[1] We define any collection of per-consumer consumptions $\boldsymbol{x} \doteq (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n) \in \mathcal{X}$ a **consumption profile**, with $\mathcal{X} \doteq \bigtimes_{i \in [n]} \mathcal{X}_i$ defined as the **set of consumption profiles**, and any collection of per-consumer endowments $\boldsymbol{e} \doteq (\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n) \in \mathbb{R}^{nm}$, as an **endowment profile**.

Remark 10.1.1.

For ease of exposition, we restrict our focus to Arrow-Debreu exchange economies, opting not to present Arrow-Debreu competitive economies (see Arrow and Debreu (1954)), which in addition to consumers are also inhabited by firms. Nevertheless, this focus on Arrow-Debreu exchange economies is without loss of generality, since any firm can be represented as a consumer in an Arrow-Debreu exchange economy by injecting an additional commodity into the economy that represents ownership of the firm, setting the

[1] In line with the literature (see, for instance, Debreu et al. (1954)), the range of this utility function should not be interpreted to have any meaning. Instead, the utility function u_i should be understood to represent a preference relation $\succeq_i$ on the space of consumptions $\mathcal{X}_i$ so that for any two consumptions $\boldsymbol{x}_i, \boldsymbol{x}'_i \in \mathcal{X}$, $u_i(\boldsymbol{x}_i) \geq u_i(\boldsymbol{x}'_i)$ implies $\boldsymbol{x}_i \succeq_i \boldsymbol{x}'_i$.

consumption space of the new consumer to be equal to the production space of the firm, and its utility function so that it seeks to maximize its consumption of the commodity associated with the firm's ownership. This commodity should also appear in the endowments of consumers that have contractual claims over the profits of the firms. A similar, albeit much more complicated reduction than described here, was proposed earlier by Garg and Kannan (2015), to which we refer the reader for additional details.

An Arrow-Debreu economy $(n, m, \mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ can be represented as a Walrasian economy $(m, \mathcal{Z})$, where the the excess demand correspondence $\mathcal{Z} : \Delta_m \rightrightarrows \mathbb{R}^m$ is given by:

$$\mathcal{Z}(\boldsymbol{p}) \doteq \sum_{i \in [n]} \left[\mathop{\arg\max}_{\boldsymbol{x}_i \in \mathcal{X}_i : \boldsymbol{x}_i \cdot \boldsymbol{p} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}} u_i(\boldsymbol{x}_i) \right] - \sum_{i \in [n]} \boldsymbol{e}_i \tag{10.1}$$

With this equivalence in hand, we define equilibrium in Arrow-Debreu economies.

Definition 10.1.1.

Given $\varepsilon \geq 0$, an $\varepsilon-$**Arrow-Debreu equilibrium** $(\boldsymbol{x}^*, \boldsymbol{p}^*)$ is a tuple comprising consumptions $\boldsymbol{x}^* \in \mathbb{R}_+^{n \times n}$ and prices $\boldsymbol{p}^* \in \Delta_m$ s.t.

(Utility maximization)	all consumers $i \in [n]$ $\varepsilon-$maximize their utility constrained by the value of their endowment: $\max_{\boldsymbol{x}_i \in \mathcal{X}_i : \boldsymbol{x}_i \cdot \boldsymbol{p}^* \leq \boldsymbol{e}_i \cdot \boldsymbol{p}^*} u_i(\boldsymbol{x}_i) \leq u_i(\boldsymbol{x}_i^*) - \varepsilon$
(Feasibility)	the consumptions are $\varepsilon-$feasible, i.e., $\sum_{i \in [n]} \boldsymbol{x}_i^* \leq \sum_{i \in [n]} \boldsymbol{e}_i - \varepsilon$
(Walras' law)	the value of the demand and the supply are equal, i.e., $\boldsymbol{p}^* \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i^* - \sum_{i \in [n]} \boldsymbol{e}_i \right) = 0$[2]

A 0-Arrow-Debreu equilibrium is simply called an Arrow-Debreu equilibrium.

Remark 10.1.2 [Arrow-Debreu equilibrium prices are Walrasian equilibria]**.**

The set of Arrow-Debreu equilibrium prices of $(n, m, \mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ is equal to the set of Walrasian equilibria of the corresponding Walrasian economy $(m, \mathcal{Z})$. To see this, notice that for any Arrow-Debreu equilibrium $(\boldsymbol{x}^*, \boldsymbol{p}^*)$ of $(n, m, \mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, we have $\sum_{i \in [n]} \boldsymbol{x}_i^* \in \mathcal{Z}(\boldsymbol{p}^*)$. Hence,

[2] We note that for the results in this section, the outputs of the algorithms we develop are guaranteed to always satisfy Walras' law, and as such, we need not introduce an ε-Walras' law.

since $\sum_{i \in [n]} \boldsymbol{x}_i^*$ is feasible and satisfies Walras' law under $\boldsymbol{p}^*$, $\boldsymbol{p}^*$ must be a Walrasian equilibrium of the Walrasian economy $(m, \mathcal{Z})$.

10.2 Solution Concepts and Existence

Remark 10.1.2 tells us that we can obtain the existence of Arrow-Debreu equilibrium prices as a corollary of the existence of Walrasian equilibria, which in turn implies the existence of Arrow-Debreu equilibrium consumptions, since for any fixed price the Arrow-Debreu equilibrium consumptions can under very mild assumption be shown to exist.

Nevertheless, a more economically meaningful proof existence can be obtained by leveraging a fundamental relationship, first observed in the seminal work of Arrow and Debreu (1954), between pseudo-games and Arrow-Debreu economies.

Definition 10.2.1 [Arrow-Debreu Pseudo-Game]**.**
Given an Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, we define the associated $(n+1)$-player **Arrow-Debreu pseudo-game** $(n+1, 1, \mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$, denoted $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$ when n is clear from context, in which the first n players are called "consumers", and the $(n+1)^{th}$ player is called the "auctioneer", and where

(Action spaces) For all consumers $i \in [n]$, $\mathcal{A}_i \doteq \mathcal{X}'_i$ and for the auctioneer $\mathcal{A}_{n+1} \doteq \Delta_m$ where, $\mathcal{X}'_i \doteq \left\{\boldsymbol{x}_i \mid \sum_{k \in [n]} \boldsymbol{x}_k \leq \sum_{k \in [n]} \boldsymbol{e}_k + \gamma, \forall \boldsymbol{x} \in \mathcal{X}\right\}$ is the **restricted consumption space** for any choice of $\gamma > 0$, which expands the set of consumptions slightly beyond those which are feasible.

(Constraints) For all consumers $i \in [n]$, $g_i(\boldsymbol{x}, \boldsymbol{p}) = \boldsymbol{p} \cdot (\boldsymbol{e}_i - \boldsymbol{x}_i)$, and for the auctioneer $g_{n+1}(\boldsymbol{x}, \boldsymbol{p}) \doteq 0$

(Payoffs) For all consumers $i \in [n]$, $u'_i(\boldsymbol{x}, \boldsymbol{p}) \doteq u_i(\boldsymbol{x}_i)$, and for the auctioneer, $u'_{n+1}(\boldsymbol{x}, \boldsymbol{p}) \doteq \boldsymbol{p} \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i - \sum_{i \in [n]} \boldsymbol{e}_i\right)$

Any Arrow-Debreu pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$ can succinctly be represented as the following $n+1$ simultaneous optimization problems:

$$\forall i \in [n], \qquad \max_{\boldsymbol{x}_i \in \mathcal{X}'_i : \boldsymbol{x}_i \cdot \boldsymbol{p} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}} u_i(\boldsymbol{x}_i) \qquad \Bigg| \qquad \max_{\boldsymbol{p} \in \Delta_m} \boldsymbol{p} \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i - \sum_{i \in [n]} \boldsymbol{e}_i \right)$$

Note that the above simultaneous $n+1$ optimization problems constitute a pseudo-game and not just a game since the prices chosen by the auctioneer determine the feasible action space of the consumers (i.e., the **budget sets** of the consumers).

With this pseudo-game in hand, next we will show that the set of Arrow-Debreu equilibria of any Arrow-Debreu economy is equal to the set of GNE of this pseudo-game. The results in this chapter will hold for (quasi)concave Arrow-Debreu economies.

Definition 10.2.2 [(Quasi)concave economies]**.**
An Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ is said to be **quasiconcave (respectively, concave)** iff it satisfies the following conditions for all consumers $i \in [n]$:

(Closed consumption set)	$\mathcal{X}_i$ is non-empty, bounded from below, closed, and convex
(Feasible budget set)	There exists a consumption that is strictly less than the consumer's endowment, i.e., for all $i \in [n]$, there exists $\boldsymbol{x}_i \in \mathcal{X}_i$ s.t. $\boldsymbol{x}_i < \boldsymbol{e}_i$
(Continuity)	u_i is continuous
((Quasi)concavity)	u_i is quasiconcave (respectively, concave)
(Non-satiation)	u_i is non-satiated, i.e., for all $\boldsymbol{x}_i \in \mathcal{X}_i$, there exists $\boldsymbol{x}'_i \in \mathcal{X}_i$ s.t. $u_i(\boldsymbol{x}'_i) > u_i(\boldsymbol{x}'_i)$

Lemma 10.2.1 [GNE = Arrow-Debreu Equilibrium]**.**
The set of Arrow-Debreu equilibria of any quasiconcave Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ is equal to the set of GNE of the corresponding Arrow-Debreu pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$.

Proof of Lemma 10.2.1

We prove only one direction (i.e., any GNE of an Arrow-Debreu pseudo-game is an Arrow-Debreu equilibrium). The converse follows similarly. For additional details, see the proof of Theorem 1 in Arrow and Debreu (1954).

Let $(\boldsymbol{x}^*, \boldsymbol{p}^*)$ be any GNE of the Arrow-Debreu pseudo-game.

First, by summing the consumers' budget constraints, and by the definition of $\boldsymbol{p}^*$ at a GNE, we have:

$$
\begin{aligned}
0 &\geq \boldsymbol{p}^* \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i^* - \sum_{i \in [n]} \boldsymbol{e}_i \right) && \text{(Sum of consumers' budget constraints)} \\
&= \max_{\boldsymbol{p} \in \Delta_m} \boldsymbol{p} \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i^* - \sum_{i \in [n]} \boldsymbol{e}_i \right) && \text{(Definition of } \boldsymbol{p}^*\text{)} \\
&\geq \boldsymbol{j}_j \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i^* - \sum_{i \in [n]} \boldsymbol{e}_i \right) && \forall j \in [m] \\
&= \sum_{i \in [n]} x_{ij}^* - \sum_{i \in [n]} e_{ij} && \forall j \in [m]
\end{aligned}
\tag{10.2}
$$

Hence, $\boldsymbol{x}^*$ is feasible: i.e., $\boldsymbol{x}^* \in \mathcal{X}'$.

Second, suppose that some consumers' budget constraints is not binding: i.e., assume there exists some consumer $i \in [n]$ s.t.

$$
\boldsymbol{x}_i^* \cdot \boldsymbol{p}^* < \boldsymbol{e}_i \cdot \boldsymbol{p}^* \tag{10.3}
$$

Now, by non-satiation, there exists $\boldsymbol{x}_i' \in \mathcal{X}_i$ s.t. $u_i(\boldsymbol{x}_i') > u_i(\boldsymbol{x}_i^*)$. As a result, there must also exist $\lambda \in (0, 1)$ s.t. for the consumption $\boldsymbol{x}_i^\dagger \doteq \lambda \boldsymbol{x}_i' + (1 - \lambda)\boldsymbol{x}_i^*$, we have:

1. $\boldsymbol{x}_i^\dagger \in \mathcal{X}_i'$, since $\mathcal{X}_i'$ is convex and $\boldsymbol{x}_i^* \in \operatorname{int}(\mathcal{X}_i')$ by Equation (10.2)
2. $u_i(\boldsymbol{x}_i^\dagger) > u_i(\boldsymbol{x}_i^*)$, since u_i is quasiconcave
3. $\boldsymbol{x}_i^\dagger \cdot \boldsymbol{p}^* \leq \boldsymbol{e}_i \cdot \boldsymbol{p}^*$, since the function $\boldsymbol{x}_i \mapsto \boldsymbol{x}_i \cdot \boldsymbol{p}^*$ is continuous

However, this is a contradiction, since $\boldsymbol{x}_i^* \in \operatorname*{arg\,max}_{\boldsymbol{x}_i \in \mathcal{X}_i : \boldsymbol{x}_i \cdot \boldsymbol{p} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}} u_i(\boldsymbol{x}_i)$. Hence, for all consumers $i \in [n]$, we must have:

$$
\boldsymbol{x}_i^* \cdot \boldsymbol{p}^* = \boldsymbol{e}_i^* \cdot \boldsymbol{p}^* \tag{10.4}
$$

Summing Equation (10.4) over consumers $i \in [n]$, and rearranging yields:

$$\boldsymbol{p}^* \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i^* - \sum_{i \in [n]} \boldsymbol{e}_i \right) = 0$$

Hence, $(\boldsymbol{x}^*, \boldsymbol{p}^*)$ satisfies Walras' law.

Finally, observe that $\boldsymbol{x}_i^* \in \text{int}(\mathcal{X}_i')$ for all consumers $i \in [n]$ by Equation (10.2). Furthermore, since $\mathcal{X}_i' \subseteq \mathcal{X}_i$, it follows that $u_i(\boldsymbol{x}_i^*) = \max_{\boldsymbol{x}_i \in \mathcal{X}_i' : \boldsymbol{x}_i \cdot \boldsymbol{p}^* \leq \boldsymbol{e}_i \cdot \boldsymbol{p}^*} u_i(\boldsymbol{x}_i) = \max_{\boldsymbol{x}_i \in \mathcal{X}_i : \boldsymbol{x}_i \cdot \boldsymbol{p}^* \leq \boldsymbol{e}_i \cdot \boldsymbol{p}^*} u_i(\boldsymbol{x}_i)$. Hence, consumers are utility maximizing constrained by the value of their endowments at prices $\boldsymbol{p}^*$, i.e., for all consumers $i \in [n]$, we have $\boldsymbol{x}_i^* \in \arg\max_{\boldsymbol{x}_i \in \mathcal{X}_i : \boldsymbol{x}_i \cdot \boldsymbol{p}^* \leq \boldsymbol{e}_i \cdot \boldsymbol{p}^*} u_i(\boldsymbol{x}_i)$.

Putting it all together, $(\boldsymbol{x}^*, \boldsymbol{p}^*)$ is then an Arrow-Debreu equilibrium.

Remark 10.2.1 [Bounded Excess Demand]**.**

Recall that in Remark 5.4.8 we had suggested that assuming boundedness of the excess demand for the convergence of mirror extratâtonnement (Algorithm 6) was a natural assumption. The above lemma provides a justification for this earlier remark since it means that in the definition of the excess demand for Arrow-Debreu markets we can replace $\mathcal{X}$ by $\mathcal{X}'$. Unlike $\mathcal{X}$, $\mathcal{X}'$ is compact. Therefore, by Berge's maximum theorem (Berge, 1997), the excess demand $\boldsymbol{z}$ is continuous over Δ_m,[3] and as such the excess demand is guaranteed to be bounded by $\max_{\boldsymbol{p} \in \Delta_m} \|\boldsymbol{z}(\boldsymbol{p})\|$, where the maximum is well-defined, since Δ_m is non-empty and compact, and since $\boldsymbol{z}$ is continuous.

With the above lemma in hand, we can apply Theorem 9.2.1 to prove the existence of an Arrow-Debreu equilibrium.

Theorem 10.2.1 [Existence of Arrow-Debreu Equilibrium]**.**

An Arrow-Debreu equilibrium is guaranteed to exist in any quasiconcave Arrow-Debreu economy.

[3]Recall that for our algorithmic results we assume $\mathcal{Z}$ is singleton-valued.

Proof of Theorem 10.2.1

Given any concave Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, construct the associated Arrow-Debreu pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$ as in Definition 10.2.1. First, notice that the action spaces of the players $\{\mathcal{X}_i'\}_i$ and Δ_m are non-empty, compact, and convex. Second, $\boldsymbol{g}$ is continuous; for all players $i \in [n+1]$, $\boldsymbol{g}_i$ is quasiconcave in the action of the i^{th} player's action; and for all players $i \in [n]$, and $\boldsymbol{p} \in \Delta_m$ there exists $\boldsymbol{x}_i$ s.t. $\boldsymbol{g}_i(\boldsymbol{x}, \boldsymbol{p}) \geq 0$. Finally, for all players $i \in [n+1]$, u_i' is continuous, as well as quasiconcave in each player's action. Hence, by Theorem 9.2.1 a GNE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$ is guaranteed to exist. In turn, by Remark 10.2.1 an Arrow-Debreu equilibrium is guaranteed to exist the Arrow-Debreu economy.

10.3 Computation of Arrow-Debreu Equilibrium

With the question of existence of Arrow-Debreu equilibria out of the way, we now turn our attention to the computation of an Arrow-Debreu equilibrium.

To motivate some of the issues that arise in the computation of an Arrow-Debreu equilibrium, we will first apply the results derived in Chapter 5 to establish the polynomial-time convergence of the mirror *extratâtonnment* process to an Arrow-Debreu equilibrium in Arrow-Debreu economies that satisfy the Minty condition assuming access to an excess demand oracle. However, as the implementation of an excess demand oracle is in general not possible, and because the mirror *extratâtonnment* process is not necessarily guaranteed to converge with only an approximate excess demand oracle, this polynomial-time computation of Walrasian equilibrium in Walrasian economies cannot be interpreted as a polynomial-time computation of an Arrow-Debreu equilibrium in Arrow-Debreu economies. We will thus introduce a new market dynamic, which we call the mirror extratrade dynamic, and we will show that this dynamic converges in polynomial time to an Arrow-Debreu equilibrium in a large class of Arrow-Debreu economies, known as pure exchange economies.

10.3.1 Computational Model

With the question of existence answered, we now turn our attention to the computation of Arrow-Debreu equilibrium.

Algorithms for the computation of an Arrow-Debreu equilibrium are called **market dynamics**.

Definition 10.3.1 [Market Dynamics]**.**
Given an Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ and an initial iterate $(\boldsymbol{p}^{(0)}, \boldsymbol{x}^{(0)}) \in \Delta_m \times \mathcal{X}$, a **market dynamic $\boldsymbol{\pi}$** consists of an update function that generates the sequence of iterates

$\{\boldsymbol{p}^{(t)}, \boldsymbol{x}^{(t)}\}_t$ given by: for all $t = 0, 1, \ldots$,

$$(\boldsymbol{p}^{(t+1)}, \boldsymbol{x}^{(t+1)}) \doteq \boldsymbol{\pi}\left((\mathcal{X}, \boldsymbol{e}, \boldsymbol{u}) \cup \bigcup_{i=0}^{t} (\boldsymbol{p}^{(i)}, \boldsymbol{x}^{(i)})\right)$$

The computational complexity results in this chapter will rely on the following computational model.

Definition 10.3.2 [Arrow-Debreu Economy Computational Model]**.**
Given an Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ s.t. for some $k \in \mathbb{N}_{++}$ the derivatives $\{\nabla^j \boldsymbol{u}\}_{j=1}^{k-1}$ are well-defined, the computational complexity of a market dynamic is measured in term of the number of evaluations of the the functions $\boldsymbol{u}, \nabla \boldsymbol{u}, \ldots, \nabla^k \boldsymbol{u}$.

10.3.2 *Tâtonnement* in WARP Arrow-Debreu Economies

Recall that any Arrow-Debreu economy $(n, m, \mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ can be represented as a Walrasian economy $(m, \mathcal{Z})$. This equivalence provides us with a first approach to computing an Arrow-Debreu equilibrium, namely computing a Walrasian equilibrium $\boldsymbol{p}^* \in \Delta_m$ of the Walrasian economy $(m, \mathcal{Z})$ using the mirror *extratâtonnement* process introduced in Chapter 5, and then setting for all consumers $i \in [n]$, $\boldsymbol{x}_i^* \dot{\in} \operatorname*{arg\,max}_{\boldsymbol{x}_i \in \mathcal{X}_i : \boldsymbol{x}_i \cdot \boldsymbol{p}^* \leq \boldsymbol{e}_i \cdot \boldsymbol{p}^*} u_i(\boldsymbol{x}_i)$, so that $(\boldsymbol{x}^*, \boldsymbol{p}^*)$ is an Arrow-Debreu equilibrium of the Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$.

Now, if we assume, that for all players $i \in [n]$, u_i is strictly concave, we can ensure that $\mathcal{Z} = \{\boldsymbol{z}\}$ is singleton-valued, and under suitable additional conditions, we can further ensure that $\mathcal{Z}$ satisfies variational stability. However, recall that in Chapter 5, to prove the convergence of the mirror *extratâtonnement* process, we had to ensure that $\boldsymbol{z}$ is bounded, which requires highly restrictive assumptions on $\boldsymbol{z}$. Nevertheless, Lemma 10.2.1 suggests that we can consider an alternative excess demand correspondence definition $\mathcal{Z}'$ for Arrow-Debreu economies s.t. $\mathcal{Z}'$ is bounded-valued, which defines a Walrasian Arrow-Debreu economy we call a **restricted Walrasian Arrow-Debreu economy**.

Definition 10.3.3 [Restricted Walrasian Arrow-Debreu Economy]**.**
Given an Arrow-Debreu economy $(n, m, \mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, the **Walrasian Arrow-Debreu economy**

$(m, \mathcal{Z})$ is a Walrasian economy with the excess demand correspondence given by:

$$\mathcal{Z}'(\boldsymbol{p}) = \sum_{i \in [n]} \left[\underset{\boldsymbol{x}_i \in \mathcal{X}'_i : \boldsymbol{x}_i \cdot \boldsymbol{p} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}}{\arg\max} u_i(\boldsymbol{x}_i) \right] - \sum_{i \in [n]} \boldsymbol{e}_i \enspace ,$$

where $\mathcal{X}'_i = \left\{ \boldsymbol{x}_i \mid \sum_{k \in [n]} \boldsymbol{x}_k \leq \sum_{k \in [n]} \boldsymbol{e}_k + \gamma, \forall k \in [n], \boldsymbol{x}_k \in \mathcal{X}_k \right\}$ for any choice of $\gamma > 0$.

Notice that in the definition of $\mathcal{Z}'$, the consumption sets $\{\mathcal{X}_i\}_i$ in the definition of $\mathcal{Z}$ (Equation (10.1)) are replaced by the restricted consumption sets $\{\mathcal{X}'_i\}_i$. From Lemma 10.2.1, we can then infer that any Walrasian equilibrium $\boldsymbol{p}^* \in \Delta_m$ of the restricted Walrasian Arrow-Debreu economy $(m, \mathcal{Z}')$ is an Arrow-Debreu equilibrium price of $(n, m, \mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$. Further, as shown in the following lemma, it is straightforward to verify that the Walrasian economy $(m, \mathcal{Z})$, as the name suggests, gives rise to a Walrasian competitive economy.

Lemma 10.3.1 [Arrow-Debreu Economies are Walrasian competitive Economies]**.** Consider the Walrasian Arrow-Debreu competitive economy $(m, \mathcal{Z})$ associated with the quasiconcave Arrow-Debreu economy $(n, m, \mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$. Then, $\mathcal{Z}$ satisfies the following:

1. (Homogeneity of degree 0) For all $\lambda > 0$, $\mathcal{Z}(\lambda \boldsymbol{p}) = \mathcal{Z}(\boldsymbol{p})$

2. (Weak Walras' law) For all $\boldsymbol{p} \in \mathbb{R}^m_+$ and $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$, $\boldsymbol{p} \cdot \boldsymbol{z}(\boldsymbol{p}) \leq 0$

3. (Non-Satiation) for all $\boldsymbol{p} \in \mathbb{R}^m_+$ and $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$, $\boldsymbol{z}(\boldsymbol{p}) \leq \boldsymbol{0}_m$ implies $\boldsymbol{p} \cdot \boldsymbol{z}(\boldsymbol{p}) = 0$

4. (Continuity) The excess demand correspondence $\mathcal{Z}$ is upper hemicontinuous on Δ_m, non-empty-, compact-, and convex-valued

5. (Boundedness) For all $\boldsymbol{p} \in \mathbb{R}^m_+$ and $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$, $\|\boldsymbol{z}(\boldsymbol{p})\|_\infty < \infty$

That is, the Walrasian Arrow-Debreu competitive economy $(m, \mathcal{Z})$ associated with the Arrow-Debreu economy $(n, m, \mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, is a continuous competitive economy, which, in addition, is bounded.

Proof of Lemma 10.3.1

Homogeneity. For all $\lambda > 0$, we have:

$$\begin{aligned}
\mathcal{Z}(\lambda \boldsymbol{p}) &= \sum_{i \in [n]} \left[\mathop{\arg\max}_{\boldsymbol{x}_i \in \mathcal{X}_i' : \boldsymbol{x}_i \cdot (\lambda \boldsymbol{p}) \leq \boldsymbol{e}_i \cdot (\lambda \boldsymbol{p})} u_i(\boldsymbol{x}_i) \right] - \sum_{i \in [n]} \boldsymbol{e}_i \\
&= \sum_{i \in [n]} \left[\mathop{\arg\max}_{\boldsymbol{x}_i \in \mathcal{X}_i' : \lambda \boldsymbol{x}_i \cdot \boldsymbol{p} \leq \lambda \boldsymbol{e}_i \cdot \boldsymbol{p}} u_i(\boldsymbol{x}_i) \right] - \sum_{i \in [n]} \boldsymbol{e}_i \\
&= \sum_{i \in [n]} \left[\mathop{\arg\max}_{\boldsymbol{x}_i \in \mathcal{X}_i' : \boldsymbol{x}_i \cdot \boldsymbol{p} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}} u_i(\boldsymbol{x}_i) \right] - \sum_{i \in [n]} \boldsymbol{e}_i \\
&= \mathcal{Z}(\boldsymbol{p})
\end{aligned}$$

Walras' law. Fix any $\boldsymbol{p} \in \mathbb{R}_+^m$. For all consumers $i \in [n]$, if $\boldsymbol{x}_i^* \dot{\in} \mathop{\arg\max}_{\boldsymbol{x}_i \in \mathcal{X}_i' : \boldsymbol{x}_i \cdot \boldsymbol{p} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}} u_i(\boldsymbol{x}_i)$, then

$$\boldsymbol{x}_i^* \cdot \boldsymbol{p} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}$$

Summing over all consumers, and rearranging yields:

$$\boldsymbol{p} \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i^* - \sum_{i \in [n]} \boldsymbol{e}_i \right) \leq 0$$

Hence, for all $\boldsymbol{p} \in \mathbb{R}_+^m$ and $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$, $\boldsymbol{p} \cdot \boldsymbol{z}(\boldsymbol{p}) \leq 0$.

Non-Satiation Fix any $\boldsymbol{p} \in \Delta_m$. For all consumers $i \in [n]$, let $\boldsymbol{x}_i^* \dot{\in} \mathop{\arg\max}_{\boldsymbol{x}_i \in \mathcal{X}_i' : \boldsymbol{x}_i \cdot \boldsymbol{p} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}} u_i(\boldsymbol{x}_i)$. Suppose not: i.e., assume $\boldsymbol{z}(\boldsymbol{p}) \leq \boldsymbol{0}_m$, while there exists some consumer $i \in [n]$ s.t.

$$\boldsymbol{x}_i^* \cdot \boldsymbol{p} < \boldsymbol{e}_i \cdot \boldsymbol{p}$$

Now, by non-satiation, there exists $\boldsymbol{x}_i' \in \mathcal{X}_i$ s.t. $u_i(\boldsymbol{x}_i') > u_i(\boldsymbol{x}_i^*)$. As a result, there must also exist $\lambda \in (0, 1)$ s.t. for the consumption $\boldsymbol{x}_i^\dagger \doteq \lambda \boldsymbol{x}_i' + (1 - \lambda)\boldsymbol{x}_i^*$,

1. $\boldsymbol{x}_i^\dagger \in \mathcal{X}_i'$ since $\boldsymbol{x}_i^* \in \operatorname{int}(\mathcal{X}_i')$
2. $u_i(\boldsymbol{x}_i^\dagger) > u_i(\boldsymbol{x}_i^*)$ since u_i is quasiconcave
3. $\boldsymbol{x}_i^\dagger \cdot \boldsymbol{p} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}$ since the function $\boldsymbol{x}_i \mapsto \boldsymbol{x}_i \cdot \boldsymbol{p}$ is continuous.

However, this is a contradiction, because $\boldsymbol{x}_i^* \in \mathop{\arg\max}_{\boldsymbol{x}_i \in \mathcal{X}_i' : \boldsymbol{x}_i \cdot \boldsymbol{p} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}} u_i(\boldsymbol{x}_i)$.

Hence, for all consumers $i \in [n]$, we must have:

$$\boldsymbol{x}_i^* \cdot \boldsymbol{p} = \boldsymbol{e}_i^* \cdot \boldsymbol{p} \tag{10.5}$$

Summing over all consumers $i \in [n]$, and rearranging yields: for all $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$,

$$0 = \boldsymbol{p}^* \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i^* - \sum_{i \in [n]} \boldsymbol{e}_i \right) = \boldsymbol{p}^* \cdot \boldsymbol{z}(\boldsymbol{p})$$

Continuity.

Since $\mathcal{X}'$ is non-empty, compact, and convex, and for all consumers $i \in [n]$, u_i is continuous and quasiconcave, and there exists $\boldsymbol{x}_i \in \mathcal{X}_i$ s.t. $\boldsymbol{x}_i < \boldsymbol{e}_i$, the assumptions of Berge's maximum theorem (Berge, 1997) hold, and the excess demand $\mathcal{Z}$ is thus upper hemicontinuous, non-empty, compact, and convex-valued over Δ_m.

Boundedness Since for all consumers $i \in [n]$, $\mathcal{X}_i$ is bounded from below, $\mathcal{X}_i'$ must also be bounded, as it is bounded from above by $\sum_{i \in [n]} \boldsymbol{e}_i$. Hence, for all consumers $i \in [n]$, $\mathcal{X}_i'$ is compact. Hence, for all $\boldsymbol{p} \in \mathbb{R}_+^m$ and $\boldsymbol{z}(\boldsymbol{p}) \in \mathcal{Z}(\boldsymbol{p})$, $\|\boldsymbol{z}(\boldsymbol{p})\|_\infty < \operatorname{diam}(\mathcal{X}_i')$.

With the above lemma in hand, we could apply Theorem 5.4.1 to show convergence of the mirror *extratâtonnement* process under the conditions derived in Chapter 5. However, such a result would rely on the existence of an exact excess demand oracle, which cannot be guaranteed to exist in general. As such, to obtain a truly polynomial-time market dynamic for Arrow-Debreu economies, we instead resort to solving not the restricted Walrasian Arrow-Debreu economy but rather the Arrow-Debreu pseudo-game.

10.3.3 Mirror Extratrade Dynamics in Pure Exchange Economies

Unfortunately, while the Arrow-Debreu pseudo-game allows us to establish the existence of an Arrow-Debreu equilibrium in Arrow-Debreu economies, as this pseudo-game does not have jointly convex constraints, we cannot apply any of our convergence results to this pseudo-game. To get around this difficulty, we will restrict our attention to pure exchange economies, and introduce an alternative formulation of Arrow-Debreu economies as a jointly convex pseudo-game.

Definition 10.3.4 [Pure Exchange Economy]**.**

A **pure exchange economy** is an Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ s.t. for all consumers $i \in [n]$,

(No debts payable) Endowments are positive, i.e., $\boldsymbol{e}_i \in \mathbb{R}^m_+$

(No commodity creation) Consumers cannot create any commodity, i.e., $\mathcal{X}_i \subseteq \mathbb{R}^m_+$

Intuitively, pure exchange economies are those Arrow-Debreu economies in which consumers 1. do not owe any amount of any commodity to any other consumer, and 2. cannot create additional commodities. We make the following observation about pure exchange economies.

Define $\oslash$ as the Hadamard division operator, i.e., $\boldsymbol{a} \oslash \boldsymbol{b} \doteq (a_i/b_i)_i$.

Remark 10.3.1 [Positive supply of commodities]**.**

Note that in quasiconcave pure exchange economies, every commodity has a strictly positive endowment, i.e, for all commodities $j \in [m]$, $\sum_{k \in [n]} e_{kj} > 0$, and as such Hadamard divisions by $\sum_{k \in [n]} \boldsymbol{e}_k$ are always well-defined. To see this, recall that in Arrow-Debreu economies, for all consumers $i \in [n]$, there exists $\boldsymbol{x}_i \in \mathcal{X}_i$ s.t. for all commodities $j \in [m]$, $x_{ij} < e_{ij}$. However, since in pure exchange economies $\mathcal{X}_i \subseteq \mathbb{R}^m_+$, this condition is only guaranteed to hold in pure exchange economies iff for all consumers $i \in [n]$ and commodities $j \in [m]$, $e_{ij} > 0$, .

In light of the above remark, we make the following simplifying assumption, without loss of generality which will lighten our notation going forward.

Assumption 10.3.1 [Normalized Aggregate Supply]**.**

Without loss of generality, for any pure exchange economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, we assume every commodity has unit aggregate supply, i.e., $\sum_{i \in [n]} \boldsymbol{e}_i = \mathbf{1}_m$.

The following remark explains why this assumption is without loss of generality.

Remark 10.3.2 [Unit Aggregate Supply]**.**

The assumption that every commodity has unit aggregate supply is without loss of generality, because commodities are divisible, and as such any pure exchange economy without aggregate unit supply $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ can be converted into a pure exchange economy $(\mathcal{X}, \boldsymbol{e}', \boldsymbol{u}')$ with aggregate unit supply, where for all consumers $i \in [n]$, $\boldsymbol{e}_i \doteq \boldsymbol{e}_i \oslash \sum_{k \in [n]} \boldsymbol{e}_k$ and $u_i'(\boldsymbol{x}_i) \doteq u_i \left(\boldsymbol{x}_i \oslash \sum_{k \in [n]} \boldsymbol{e}_k \right)$. Then, any Arrow-Debreu equilibrium $(\boldsymbol{x}^*, \boldsymbol{p}^*)$ of $(\mathcal{X}, \boldsymbol{e}', \boldsymbol{u}')$ can be converted to an Arrow-Debreu equilibrium $(\boldsymbol{x}^{**}, \boldsymbol{p}^*)$ of $(\mathcal{X}, \boldsymbol{e}', \boldsymbol{u}')$ by simply setting, for all consumers $i \in [n]$ and commodities $j \in [m]$, $x_{ij}^{**} \doteq x_{ij}^*(\sum_i e_{ij}')$. A straightforward algebraic manipulation then verifies this constructed Arrow-Debreu equilibrium satisfies feasibility, Walras' law, and utility maximization.

Intuitively, this construction tells us that Arrow-Debreu equilibrium consumptions of any pure exchange economy with unit aggregate supply can be interpreted as equilibrium percentages of the aggregate supply consumed in any pure exchange economy without aggregate unit supply.

We now introduce an alternative pseudo-game formulation of pure exchange economies, which has jointly convex constraints. The pseudo-game we propose is equivalent to the Trading Post game (Shapley and Shubik, 1977) proposed by Shapley and Shubik, up to a variable substitution.

Definition 10.3.5 [Trading Post Pseudo-Game]**.**

Given a pure exchange economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, we define the associated n-player **trading post pseudo-game** $(n, m+3, \mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$, denoted $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$ when n and m are clear from context, in which the players are called consumers, and:

(Action spaces) For all consumers $i \in [n]$, $\mathcal{A}_i \doteq \mathcal{B}_i = \left\{ (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \in \mathbb{R}^m \times \mathbb{R}^m_{++} \mid (\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) \in \mathcal{X}_i' \right\}$, where $\mathcal{X}_i'$ is the restricted consumption set as defined in Definition 10.2.1.

(Constraints) For all consumers $i \in [n]$, $g_{i1}(\boldsymbol{\beta}, \boldsymbol{\pi}) = \boldsymbol{e}_i \cdot \left(\sum_{k \in [n]} \boldsymbol{\beta}_k\right) - \sum_{j \in [m]} \beta_{ij}$, $g_{i2}(\boldsymbol{\beta}, \boldsymbol{\pi}) = \sum_{i \in [n], j \in [m]} \beta_{ij} - 1$, $g_{i3}(\boldsymbol{\beta}, \boldsymbol{\pi}) = 1 - \sum_{i \in [n], j \in [m]} \beta_{ij}$, and $g_{i(j+3)}(\boldsymbol{\beta}, \boldsymbol{\pi}) \doteq \pi_{ij} - \sum_{k \in [n]} \beta_{kj}$, for all $j \in [m]$.

(Payoffs) For all consumers $i \in [n]$, $u'_i(\boldsymbol{\beta}, \boldsymbol{\pi}) \doteq u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i)$.

The trading post pseudo-game can succinctly be represented as the following n simultaneous optimization problems: for all consumers $i \in [n]$,

$$
\begin{aligned}
\max_{(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \in \mathcal{B}_i} \quad & u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) & \\
\text{s.t.} \quad & \sum_{j \in [m]} \beta_{ij} = \boldsymbol{e}_i \cdot \left(\sum_{k \in [n]} \boldsymbol{\beta}_k\right) & \text{(Budget constraint)} \\
\forall j \in [m], \quad & \sum_{k \in [n]} \beta_{kj} \leq \pi_{ij} & \text{(Bid constraint)} \\
& \sum_{i \in [n], j \in [m]} \beta_{ij} = 1 & \text{(Bid constraint)}
\end{aligned}
$$

With this definition in hand, we first show that set of Arrow-Debreu equilibria of any quasiconcave (Arrow-Debreu) pure exchange economy can be converted to a GNE of the associated trading post game, and vice-versa.

Lemma 10.3.2 [Trading Post GNE define Arrow-Debreu Equilibria]**.**

Consider a quasiconcave pure exchange economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ and the associated trading post pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$. If $(\beta^*, \boldsymbol{p}^*)$ is a GNE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$, then $(\boldsymbol{x}^*, \boldsymbol{p}^*)$ is an Arrow-Debreu equilibrium of $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, where

$$
\boldsymbol{p}^* \doteq \sum_{k \in [n]} \boldsymbol{\beta}_k^* \qquad \forall i \in [n], \quad \boldsymbol{x}_i^* \doteq \boldsymbol{\beta}_i^* \oslash \boldsymbol{\pi}_i^*
$$

Proof of Lemma 10.3.2

Let $(\boldsymbol{\beta}^*, \boldsymbol{\pi}^*)$ be a GNE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$. Define $\boldsymbol{x}^*$ and $\boldsymbol{p}^*$ as follows:

$$
\boldsymbol{p}^* \doteq \sum_{k \in [n]} \boldsymbol{\beta}_k^* \qquad \forall i \in [n], \quad \boldsymbol{x}_i \doteq \boldsymbol{\beta}_i^* \oslash \boldsymbol{\pi}_i^*
$$

First, note that $\boldsymbol{x}_i = \boldsymbol{\beta}_i^* \oslash \boldsymbol{\pi}_i^* \in \mathcal{X}_i$ and $\boldsymbol{p}^* \in \Delta_m$, since $\sum_{j \in [m]} p_i^* = \sum_{k \in [n], j \in [m]} \beta_{kj}^* = 1$.

Second, summing $\boldsymbol{x}_i^*$ over all consumers $i \in [n]$ and using the bid constraint $(\sum_{k\in[n]} \boldsymbol{\beta}_k) \oslash (\sum_{k\in[n]} \boldsymbol{e}_k) \leq \boldsymbol{\pi}_i$, yields: for all $j \in [m]$,

$$\sum_{i\in[n]} x_{ij}^* = \sum_{i\in[n]} \frac{\beta_{ij}^*}{\pi_{ij}^*} \leq \sum_{i\in[n]} \frac{\beta_{ij}^*}{\sum_{k\in[n]} \beta_{kj}^*} = \frac{\sum_{i\in[n]} \beta_{ij}^*}{\sum_{k\in[n]} \beta_{kj}^*} = \mathbf{1}_m = \sum_{i\in[n]} e_{ij} \tag{10.6}$$

Hence, $\boldsymbol{x}^* \leq \sum_{i\in[n]} \boldsymbol{e}_i$, meaning $\boldsymbol{x}^*$ is feasible.

Third, suppose not: i.e., assume a consumer $i \in [n]$ and a commodity $j \in [m]$ s.t. $p_j^* > 0$ and $\sum_{k\in[n]} \beta_{kj}^* < \pi_{ij}^*$. Recall that, by non-satiation, there exists $\boldsymbol{x}_i' \in \mathcal{X}_i$ s.t. $u_i(\boldsymbol{x}_i') > u_i(\boldsymbol{x}_i^*)$. Choose $\boldsymbol{\pi}_i' \in \mathbb{R}_{++}^m$ s.t. $\boldsymbol{x}_i' = \boldsymbol{\beta}_i^* \oslash \boldsymbol{\pi}_i'$. Now, we can choose a small enough $\lambda \in (0,1)$ s.t. for $\boldsymbol{x}_i^\dagger \doteq \mu \boldsymbol{x}_i' + (1-\mu)\boldsymbol{x}_i^*$, we have:

1. $\boldsymbol{x}_i^\dagger \doteq \boldsymbol{\beta}_i^* \oslash \boldsymbol{\pi}_i^\dagger$, for some $\boldsymbol{\pi}_i^\dagger$ s.t. $(\boldsymbol{\beta}_i^*, \boldsymbol{\pi}_i^\dagger) \in \mathcal{B}_i$, since $(\boldsymbol{\beta}_i^*, \boldsymbol{\pi}_i^*) \in \operatorname{int}(\mathcal{B}_i)$ by Equation (10.6)
2. $\sum_{k\in[n]} \beta_{kj}^* \leq \pi_{ij}^\dagger$, since both sides of the inequality are continuous in $\boldsymbol{\pi}$

Further, observe that $u_i(\boldsymbol{\beta}_i^* \oslash \boldsymbol{\pi}_i^\dagger) = u_i(\boldsymbol{x}_i^\dagger) > u_i(\boldsymbol{x}_i^*) = u_i(\boldsymbol{\beta}_i^* \oslash \boldsymbol{\pi}_i^*)$, since u_i is quasiconcave. However, this contradicts, the claim that $(\boldsymbol{\beta}^*, \boldsymbol{\pi}^*)$ is a GNE. Hence, for all consumers $i \in [n]$ and commodities $j \in [m]$ s.t. $p_j^* > 0$, it holds that $\sum_{k\in[n]} \beta_{kj}^* = \pi_{ij}^*$.

Hence, for all consumers $i \in [n]$,

$$\begin{aligned}
\sum_{j\in[m]} \beta_{ij}^* &= \boldsymbol{e}_i \cdot \left(\sum_{k\in[n]} \boldsymbol{\beta}_k^*\right) \\
\sum_{j\in[m]} x_{ij}^* \pi_{ij}^* &= (\boldsymbol{e}_i \cdot \boldsymbol{p}^*) \\
\sum_{j\in[m]:p_j^*>0} x_{ij}^* \pi_{ij}^* + \sum_{j\in[m]:p_j^*=0} \underbrace{x_{ij}^*}_{=0} \pi_{ij}^* &= \boldsymbol{e}_i \cdot \boldsymbol{p}^* \\
\sum_{j\in[m]:p_j^*>0} x_{ij}^* p_j^* + \sum_{j\in[m]:p_j^*=0} \underbrace{x_{ij}^*}_{=0} p_j^* &= \boldsymbol{e}_i \cdot \boldsymbol{p}^* \\
\boldsymbol{x}_i^* \cdot \boldsymbol{p}^* &= \boldsymbol{e}_i \cdot \boldsymbol{p}^*
\end{aligned}$$

Summing over all consumers $i \in [n]$, and rearranging yields:

$$\boldsymbol{p}^* \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i^* - \sum_{i \in [n]} \boldsymbol{e}_i \right) = 0$$

Hence, $(\boldsymbol{x}^*, \boldsymbol{p}^*)$ satisfies Walras' law.

Finally, since for all consumers $i \in [n]$, $\boldsymbol{x}_i^* \in \operatorname{int}(\mathcal{X}_i')$ by Equation (10.2), and $\mathcal{X}_i' \subseteq \mathcal{X}_i$, this implies that $u_i(\boldsymbol{x}_i^*) = \max_{\boldsymbol{x}_i \in \mathcal{X}_i' : \boldsymbol{x}_i \cdot \boldsymbol{p}^* \leq \boldsymbol{e}_i \cdot \boldsymbol{p}^*} u_i(\boldsymbol{x}_i) = \max_{\boldsymbol{x}_i \in \mathcal{X}_i : \boldsymbol{x}_i \cdot \boldsymbol{p}^* \leq \boldsymbol{e}_i \cdot \boldsymbol{p}^*} u_i(\boldsymbol{x}_i)$. Hence, consumers are utility maximizing constrained by the value of their endowments at prices $\boldsymbol{p}^*$, i.e., for all consumers $i \in [n]$, $\boldsymbol{x}_i^* \in \operatorname*{arg\,max}_{\boldsymbol{x}_i \in \mathcal{X}_i : \boldsymbol{x}_i \cdot \boldsymbol{p}^* \leq \boldsymbol{e}_i \cdot \boldsymbol{p}^*} u_i(\boldsymbol{x}_i)$.

Putting it all together, $(\boldsymbol{x}^*, \boldsymbol{p}^*)$ is an Arrow-Debreu equilibrium.

Remark 10.3.3 [Arrow-Debreu equilibria define trading post GNE]**.**

While we do not present a statement and proof for the converse of the above lemma, as we will not make use of it, using a similar argument to that provided in Lemma 10.3.2, we can also show that any Arrow-Debreu equilibrium can used to construct a trading post GNE. The converse direction is less useful at present since unlike in Lemma 10.2.1, where we have shown that the set of GNE of the Arrow-Debreu pseudo-game is equal to the set of Arrow-Debreu equilibria in the corresponding the Arrow-Debreu economy, for the trading post pseudo-game, we instead have an equivalence after suitable algebraic manipulation of the GNE of the trading post pseudo-game and the Arrow-Debreu equilibrium of a quasiconcave pure exchange economy. That is, when taken in conjuction with its converse, Lemma 10.3.2 cannot be interpreted an "if and only if" result, but perhaps rather as a polynomial-time equivalence result.

With the above lemma in hand, we now describe the properties of the trading post pseudo-game with the following result.

Lemma 10.3.3 [Trading Post Pseudo-Game Properties]**.**

The trading post pseudo-game associated with any quasiconcave pure exchange economy is a quasiconcave pseudo-game with jointly convex constraints.

Proof of Lemma 10.3.3

The trading post pseudo-game is jointly convex: First, note that for all consumers $i \in [n]$, $\mathcal{B}_i$ is non-empty, compact, convex since $\mathcal{X}'_i$ is non-empty, compact, and convex, and $\mathcal{B}_i$ is the perspective transformation of $\mathcal{X}'_i$ (see, for instance, Section 2.3.3 of Boyd et al. (2004)), which is continuous and preserves convexity.

Second, the constraints of the pseudo-game are all affine in $(\boldsymbol{\beta}, \boldsymbol{\pi})$, and are hence continuous and concave in $(\boldsymbol{\beta}, \boldsymbol{\pi})$. As such, the trading pseudo-game has jointly convex constraints.

The trading post pseudo-game is quasiconcave: Recall that in the trading pseudo-game, for all consumers $i \in [n]$, $u'_i(\boldsymbol{\beta}, \boldsymbol{\pi}) \doteq u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i)$. The function u_i is quasiconcave iff for all $\alpha \in \mathbb{R}$, the superlevel sets $\mathcal{X}_i^\alpha \doteq \{\boldsymbol{x}_i \in \mathcal{X}_i \mid u_i(\boldsymbol{x}_i) \geq \alpha\}$ are convex. Now, if we pass the superlevel sets $\mathcal{X}_i^\alpha$ through the mapping $\mathcal{X}_i^\alpha \rightrightarrows \{(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \in \mathbb{R}_+^m \times \mathbb{R}_{++}^m \mid \boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i \in \mathcal{X}_i^\alpha\}$, since the mapping is a linear-fractional transformation of $\mathcal{X}_i^\alpha$, the transformed sets are convex as well (see, for instance, Section 2.3.3 of Boyd et al. (2004)). That is, for all $\alpha \in \mathbb{R}$, the superlevel sets $\{(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \in \mathbb{R}_+^m \times \mathbb{R}_{++}^m \mid u'_i(\boldsymbol{\beta}, \boldsymbol{\pi}) \geq \alpha\} = \{(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \in \mathbb{R}_+^m \times \mathbb{R}_{++}^m \mid u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) \geq \alpha\}$ of u'_i are convex. Hence, u'_i is quasiconcave.

Using the above lemma, we can in turn obtain the existence of a VE in the trading post pseudo-game as a corollary of Lemma 10.3.3, because the trading post pseudo-game satisfies the conditions for the existence of a VE (Theorem 9.2.2).

Corollary 10.3.1.

The set of VE of the trading post pseudo-game associated with any Arrow-Debreu economy is non-empty.

Since any VE is a GNE, by Lemma 10.3.2, we can set our sights on the computation of a VE in the trading post pseudo-game to compute an Arrow-Debreu equilibrium in a pure exchange economy. Our first approach to computing a VE will be to apply the mirror extragradient learning dynamics to the trading post pseudo-game. We call the market dynamics arising from running the mirror extragradient learning dynamics on the trading post pseudo-game the **mirror extratrade dynamics**. To guarantee convergence of the mirror extragradient learning dynamics in the trading post pseudo-game, we must ensure that the trading post pseudo-game is variationally stable. It turns out that when the utility functions of the consumers in a pure exchange economy are concave (rather than quasiconcave), the associated trading post pseudo-game is guaranteed to be variationally stable. To this end, we introduce the following technical lemma.

Lemma 10.3.4 [Pseudoconcavity of composition of concave and ratio functions]**.**
If $f : \mathbb{R}^n_+ \to \mathbb{R}$ is a concave function and $\boldsymbol{g} : \mathbb{R}^n_+ \times \mathbb{R}^n_{++} \to \mathbb{R}^n_+$ is the ratio function s.t. $\boldsymbol{g}(\boldsymbol{a}, \boldsymbol{b}) \doteq \boldsymbol{a} \oslash \boldsymbol{b}$, then $\nu(\boldsymbol{a}, \boldsymbol{b}) \doteq f(\boldsymbol{g}(\boldsymbol{a}, \boldsymbol{b}))$ is pseudoconcave.

Proof of Lemma 10.3.4

Assume $\boldsymbol{a} \in \mathbb{R}^n_+$ and $\boldsymbol{b} \in \mathbb{R}^n_{++}$ s.t.

$$0 \geq \max_{\boldsymbol{a}' \in \mathbb{R}^n_+, \boldsymbol{b}' \in \mathbb{R}^n_{++}} \left\langle \nabla \nu(\boldsymbol{a}, \boldsymbol{b}), (\boldsymbol{a}', \boldsymbol{b}') - (\boldsymbol{a}, \boldsymbol{b}) \right\rangle$$

Now, since $\nabla_{a_i}\nu(\boldsymbol{a},\boldsymbol{b}) = \frac{1}{b_i}\nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b})$, and $\nabla_{b_i}\nu(\boldsymbol{a},\boldsymbol{b}) = -\frac{a_i}{b_i^2}\nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b})$, it holds that:

$$
\begin{aligned}
0 &\geq \max_{\boldsymbol{a}' \in \mathbb{R}^n_+, \boldsymbol{b}' \in \mathbb{R}^n_{++}} \left\langle \nabla \nu(\boldsymbol{a},\boldsymbol{b}), (\boldsymbol{a}',\boldsymbol{b}') - (\boldsymbol{a},\boldsymbol{b}) \right\rangle \\
&= \max_{\boldsymbol{a}' \in \mathbb{R}^n_+, \boldsymbol{b}' \in \mathbb{R}^n_{++}} \sum_{i \in [n]} \left\langle \nabla_{(a_i,b_i)} \nu(\boldsymbol{a},\boldsymbol{b}), (a_i', b_i') - (a_i, b_i) \right\rangle \\
&= \max_{\boldsymbol{a}' \in \mathbb{R}^n_+, \boldsymbol{b}' \in \mathbb{R}^n_{++}} \sum_{i \in [n]} \left\langle \left(\frac{1}{b_i} \nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b}), -\frac{a_i}{b_i^2} \nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b}) \right), (a_i', b_i') - (a_i, b_i) \right\rangle \\
&= \max_{\boldsymbol{a}' \in \mathbb{R}^n_+, \boldsymbol{b}' \in \mathbb{R}^n_{++}} \sum_{i \in [n]} \left[\frac{1}{b_i} \left\langle \nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b}), a_i' - a_i \right\rangle - \frac{a_i}{b_i^2} \left\langle \nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b}), b_i' - b_i \right\rangle \right] \\
&= \max_{\boldsymbol{a}' \in \mathbb{R}^n_+, \boldsymbol{b}' \in \mathbb{R}^n_{++}} \sum_{i \in [n]} \left[\left\langle \nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b}), \frac{a_i'}{b_i} - \frac{a_i}{b_i} \right\rangle + \left\langle \nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b}), \frac{a_i}{b_i} - \frac{a_i b_i'}{b_i^2} \right\rangle \right] \\
&= \max_{\boldsymbol{a}' \in \mathbb{R}^n_+, \boldsymbol{b}' \in \mathbb{R}^n_{++}} \sum_{i \in [n]} \left\langle \nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b}), \frac{a_i'}{b_i} - \frac{a_i b_i'}{b_i^2} \right\rangle \\
&= \max_{\boldsymbol{a}' \in \mathbb{R}^n_+, \boldsymbol{b}' \in \mathbb{R}^n_{++}} \sum_{i \in [n]} \left\langle \nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b}), \frac{b_i'}{b_i} \left(\frac{a_i'}{b_i'} - \frac{a_i}{b_i} \right) \right\rangle \\
&= \sum_{i \in [n]} \underbrace{\max_{a_i', b_i'} \left\{ \frac{b_i'}{b_i} \left\langle \nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b}), \frac{a_i'}{b_i'} - \frac{a_i}{b_i} \right\rangle \right\}}_{\geq \frac{b_i}{b_i} \left\langle \nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b}), \frac{a_i}{b_i} - \frac{a_i}{b_i} \right\rangle = 0} \\
&\geq 0
\end{aligned}
$$

Now, since for all $i \in [n]$, $\frac{b_i'}{b_i} > 0$ since $b_i, b_i' \in \mathbb{R}_{++}$, and $\max_{a_i', b_i'} \left\{ \frac{b_i'}{b_i} \left\langle \nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b}), \frac{a_i'}{b_i'} - \frac{a_i}{b_i} \right\rangle \right\} \geq 0$, for the above inequality to hold at $(\boldsymbol{a}, \boldsymbol{b})$, it must be that for all $\boldsymbol{a}' \in \mathbb{R}^n_+, \boldsymbol{b}' \in \mathbb{R}^n_{++}$ and for all $i \in [n]$,

$$
0 \geq \left\langle \nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b}), \frac{a_i'}{b_i'} - \frac{a_i}{b_i} \right\rangle \quad \text{implies} \quad 0 \geq \sum_{i \in [n]} \left\langle \nabla_i f(\boldsymbol{a} \oslash \boldsymbol{b}), \frac{a_i'}{b_i'} - \frac{a_i}{b_i} \right\rangle
$$

Hence, since f is concave, for all $\boldsymbol{a}' \in \mathbb{R}_+^n, \boldsymbol{b}' \in \mathbb{R}_{++}^n$,

$$
\begin{aligned}
0 &\geq \sum_{i \in [n]} \left\langle \nabla_i f\left(\boldsymbol{a} \oslash \boldsymbol{b}\right), \frac{a_i'}{b_i'} - \frac{a_i}{b_i} \right\rangle \\
&= \left\langle \nabla f\left(\boldsymbol{a} \oslash \boldsymbol{b}\right), \left(\boldsymbol{a}' \oslash \boldsymbol{b}'\right) - \left(\boldsymbol{a} \oslash \boldsymbol{b}\right) \right\rangle \\
&\geq f\left(\boldsymbol{a}' \oslash \boldsymbol{b}'\right)\right) - f\left(\boldsymbol{a} \oslash \boldsymbol{b}\right) \\
&= \nu(\boldsymbol{a}', \boldsymbol{b}') - \nu(\boldsymbol{a}, \boldsymbol{b})
\end{aligned}
$$

Putting it all together, for all $\boldsymbol{a}, \boldsymbol{a}' \in \mathbb{R}_+^n$ and $\boldsymbol{b}, \boldsymbol{b}' \in \mathbb{R}_{++}^n$:

$$
\left\langle \nabla \nu(\boldsymbol{a}, \boldsymbol{b}), (\boldsymbol{a}', \boldsymbol{b}') - (\boldsymbol{a}, \boldsymbol{b}) \right\rangle \implies \nu(\boldsymbol{a}, \boldsymbol{b}) \geq \nu(\boldsymbol{a}', \boldsymbol{b}')
$$

That is, ν is pseudoconcave.

Lemma 10.3.5 [Variational Stability in the Trading Post Pseudo-Game]**.**
Consider a concave pure exchange economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$. The trading post pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$ associated with $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ is pseudomonotone and variationally stable.

Proof of Lemma 10.3.5

Let $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$ be the the trading post pseudo-game associated with the pure exchange economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, where for all consumers $i \in [n]$, the utility function u_i is concave. First, for all consumers $i \in [n]$, u_i' depends only on $(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i)$, hence, we redefine $u_i'(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \doteq u_i'(\boldsymbol{\beta}, \boldsymbol{\pi})$.

Now consider the utilitarian welfare function $w(\boldsymbol{x}) \doteq \sum_{i \in [n]} u_i(\boldsymbol{x}_i)$. Since for all consumers $i \in [n]$, u_i is concave, w must also be concave, as it is the sum of concave functions (see, for instance, Section 3.2 of Boyd et al. (2004)). If we then define $w'(\boldsymbol{\beta}, \boldsymbol{\pi}) \doteq w(\boldsymbol{\beta} \oslash \boldsymbol{\pi}) = \sum_{i \in [n]} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) = \sum_{i \in [n]} u_i'(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i)$, by Lemma 10.3.4, w' must be pseudoconcave, since it is the composition of a concave function, namely w, and the ratio function. Since w' is pseudcoconcave, $-w'$ is pseudoconvex, and hence $-\nabla w'$ must be pseudomonotone (see, for instance, Theorem 4.1 of (Aussel et al., 1994)). Furthermore,

$$\nabla w'(\boldsymbol{\beta}, \boldsymbol{\pi}) = (\nabla u_1'(\boldsymbol{\beta}_1, \boldsymbol{\pi}_1), \dots, \nabla u_n'(\boldsymbol{\beta}_n, \boldsymbol{\pi}_n))$$

To prove that the trading post pseudo-game is variationally stable (Definition 9.4.2), we must show that there exists $(\boldsymbol{\beta}^*, \boldsymbol{\pi}^*) \in \mathcal{X}^*$ s.t.

$$\sum_{i \in [n]} \left\langle \nabla u_i'(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i), (\boldsymbol{\beta}_i^*, \boldsymbol{\pi}_i^*) - (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \right\rangle \geq 0 \qquad \forall (\boldsymbol{\beta}, \boldsymbol{\pi}) \in \mathcal{X}^*$$

We define the pseudo-game operator associated with the trading post pseudo-game as follows:

$$\boldsymbol{v}(\boldsymbol{\beta}, \boldsymbol{\pi}) \doteq -(\nabla u_1'(\boldsymbol{\beta}_1, \boldsymbol{\pi}_1), \dots, \nabla u_n'(\boldsymbol{\beta}_n, \boldsymbol{\pi}_n))$$

Then, equivalently, the trading post pseudo-game is variationally stable iff the set of weak solutions of the VI $(\mathcal{X}^*, \boldsymbol{v}) = (\mathcal{X}^*, -\nabla w')$ is non-empty. Since $-\nabla w'$ is pseudomonotone and $\mathcal{X}^*$ is non-empty and compact (Lemma 10.3.3), a weak solution of $(\mathcal{X}^*, -\nabla w')$ is guaranteed to exist. That is, there exists $(\boldsymbol{\beta}^*, \boldsymbol{\pi}^*) \in \mathcal{X}^*$ s.t. for all $(\boldsymbol{\beta}, \boldsymbol{\pi}) \in \mathcal{X}^*$,

$$\left\langle -\nabla w'(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i), (\boldsymbol{\beta}_i^*, \boldsymbol{\pi}_i^*) - (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \right\rangle = \sum_{i \in [n]} \left\langle \nabla u_i'(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i), (\boldsymbol{\beta}_i^*, \boldsymbol{\pi}_i^*) - (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \right\rangle \geq 0$$

With the sufficient conditions for the trading post pseudo-game to be variationally stable in hand, to guarantee polynomial-time convergence of the mirror extragradient learning dynamics, we have to ensure that the trading post pseudo-game is Lipschitz-smooth. The following lemma provides necessary conditions to ensure the Lipschitz-smoothness of the trading post pseudo-game.

Lemma 10.3.6 [Lipschitz-smoothness of Trading Post Pseudo-Game]**.**
Consider a pure exchange economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$ s.t. for all consumers $i \in [n]$, u_i is ℓ-Lipschitz-continuous, and ∇u_i is λ-Lipschitz-continuous, and also the trading post pseudo-game

$(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$ associated with the pure exchange economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$. If we define:

$$\pi_{\min} \doteq \min\{\pi_{ij} : (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \in \mathcal{B}_i, \forall i \in [n], j \in [m]\} \ ,$$

$$\pi_{\max} \doteq \max\{\pi_{ij} : (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \in \mathcal{B}_i, \forall i \in [n], j \in [m]\} \ , \text{ and}$$

$$\beta_{\max} \doteq \max\{\beta_{ij} : (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \in \mathcal{B}_i, \forall i \in [n], j \in [m]\} \ ,$$

then for all consumers $i \in [n]$, u_i' is $m \left[\max\left(\frac{\lambda}{\pi_{\min}^2}, \frac{\beta_{\max}'\lambda}{\pi_{\min}^2}, \frac{\ell}{\pi_{\min}^2}\right) + \max\left(\frac{\pi_{\max}\lambda}{\pi_{\min}}, \frac{\pi_{\max}\lambda}{\pi_{\min}^2}, \ell\right)\right]$ $-$Lipschitz-smooth.

Proof of Lemma 10.3.6

First, note that for all consumers $i \in [n]$ and commodities $j \in [m]$,

$$\begin{aligned}
\nabla_{\beta_{ij}} u_i'(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) &= \nabla_{\beta_{ij}} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) \\
&= \frac{\nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i)}{\pi_{ij}} \ ,
\end{aligned}$$

and

$$\begin{aligned}
\nabla_{\pi_{ij}} u_i'(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) &= \nabla_{\pi_{ij}} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) \\
&= -\pi_{ij}^2 \nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) \ .
\end{aligned}$$

Define $\odot$ as the Hadamard product operator, i.e., $\boldsymbol{a} \odot \boldsymbol{b} \doteq (a_i b_i)_i$. Now, for all $i \in [n]$, $j \in [m]$, $(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i), (\boldsymbol{\beta}_i', \boldsymbol{\pi}_i') \in \mathcal{B}_i$,

$$\begin{aligned}
&\left|\frac{\nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i)}{\pi_{ij}} - \frac{\nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i' \oslash \boldsymbol{\pi}_i')}{\pi_{ij}'}\right| \\
&= \left|\frac{1}{\pi_{ij}}\left(\nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) - \nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i' \oslash \boldsymbol{\pi}_i')\right) + \nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i' \oslash \boldsymbol{\pi}_i')\left(\frac{1}{\pi_{ij}} - \frac{1}{\pi_{ij}'}\right)\right| \\
&\leq \left|\frac{1}{\pi_{ij}}\left(\nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) - \nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i' \oslash \boldsymbol{\pi}_i')\right)\right| + \left|\nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i' \oslash \boldsymbol{\pi}_i')\left(\frac{1}{\pi_{ij}} - \frac{1}{\pi_{ij}'}\right)\right|
\end{aligned}$$

$$
\begin{aligned}
&= \frac{1}{\pi_{ij}} \left| \left(\nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) - \nabla_{x_{ij}} u_i(\boldsymbol{\beta}'_i \oslash \boldsymbol{\pi}'_i) \right) \right| + \left| \nabla_{x_{ij}} u_i(\boldsymbol{\beta}'_i \oslash \boldsymbol{\pi}'_i) \right| \left| \frac{1}{\pi_{ij}} - \frac{1}{\pi'_{ij}} \right| \\
&\leq \frac{\lambda}{\pi_{ij}} \left\| (\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) - (\boldsymbol{\beta}'_i \oslash \boldsymbol{\pi}'_i) \right\| + \ell \left| \frac{1}{\pi_{ij}} - \frac{1}{\pi'_{ij}} \right| \\
&= \frac{\lambda}{\pi_{ij}} \left\| (\boldsymbol{\beta}_i - \boldsymbol{\beta}'_i) \oslash \boldsymbol{\pi}_i - \boldsymbol{\beta}'_i \odot \left(\mathbf{1}_m \oslash \boldsymbol{\pi}_i - \mathbf{1}_m \oslash \boldsymbol{\pi}'_i \right) \right\| + \ell \left| \frac{1}{\pi_{ij}} - \frac{1}{\pi'_{ij}} \right| \\
&\leq \frac{\lambda}{\pi_{ij}} \left\| (\boldsymbol{\beta}_i - \boldsymbol{\beta}'_i) \oslash \boldsymbol{\pi}_i \right\| + \left\| \boldsymbol{\beta}'_i \odot \left(\mathbf{1}_m \oslash \boldsymbol{\pi}_i - \mathbf{1}_m \oslash \boldsymbol{\pi}'_i \right) \right\| + \ell \left| \frac{1}{\pi_{ij}} - \frac{1}{\pi'_{ij}} \right| \\
&\leq \frac{\lambda}{\pi_{ij} \min_{k \in [m]} \{\pi_{ik}\}} \left\| \boldsymbol{\beta}_i - \boldsymbol{\beta}'_i \right\| + \lambda \left\| \boldsymbol{\beta}'_i \right\| \left\| \mathbf{1}_m \oslash \boldsymbol{\pi}_i - \mathbf{1}_m \oslash \boldsymbol{\pi}'_i \right\| + \ell \left| \frac{1}{\pi_{ij}} - \frac{1}{\pi'_{ij}} \right|
\end{aligned}
$$

$$
\begin{aligned}
&= \frac{\lambda}{\pi_{ij} \min_{k \in [m]} \{\pi_{ik}\}} \left\| \boldsymbol{\beta}_i - \boldsymbol{\beta}'_i \right\| + \lambda \left\| \boldsymbol{\beta}'_i \right\| \left\| \left(\boldsymbol{\pi}'_i - \boldsymbol{\pi}_i \right) \oslash \left(\boldsymbol{\pi}_i \odot \boldsymbol{\pi}'_i \right) \right\| + \ell \left| \frac{\pi'_{ij} - \pi_{ij}}{\pi_{ij} \pi'_{ij}} \right| \\
&\leq \frac{\lambda}{\pi_{ij} \min_{k \in [m]} \{\pi_{ik}\}} \left\| \boldsymbol{\beta}_i - \boldsymbol{\beta}'_i \right\| + \frac{\lambda \left\| \boldsymbol{\beta}'_i \right\|}{\min_{k \in [m]} \{\pi_{ik} \pi'_{ik}\}} \left\| \boldsymbol{\pi}'_i - \boldsymbol{\pi}_i \right\| + \ell \frac{1}{\pi_{ij} \pi'_{ij}} \left| \pi'_{ij} - \pi_{ij} \right| \\
&\leq \frac{\lambda}{\pi_{ij} \min_{k \in [m]} \{\pi_{ik}\}} \left\| \boldsymbol{\beta}_i - \boldsymbol{\beta}'_i \right\| + \frac{\lambda \left\| \boldsymbol{\beta}'_i \right\|}{\min_{k \in [m]} \{\pi_{ik} \pi'_{ik}\}} \left\| \boldsymbol{\pi}'_i - \boldsymbol{\pi}_i \right\| + \ell \frac{1}{\pi_{ij} \pi'_{ij}} \left\| \boldsymbol{\pi}'_i - \boldsymbol{\pi}_i \right\| \\
&\leq \frac{\lambda}{\pi_{\min}^2} \left\| \boldsymbol{\beta}_i - \boldsymbol{\beta}'_i \right\| + \frac{\beta'_{\max} \lambda}{\pi_{\min}^2} \left\| \boldsymbol{\pi}'_i - \boldsymbol{\pi}_i \right\| + \ell \frac{1}{\pi_{\min}^2} \left\| \boldsymbol{\pi}'_i - \boldsymbol{\pi}_i \right\| \\
&\leq \max \left(\frac{\lambda}{\pi_{\min}^2}, \frac{\beta'_{\max} \lambda}{\pi_{\min}^2}, \frac{\ell}{\pi_{\min}^2} \right) \left\| (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) - (\boldsymbol{\beta}'_i, \boldsymbol{\pi}'_i) \right\|
\end{aligned}
$$

Similarly, for all $i \in [n]$, $j \in [m]$, $(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i), (\boldsymbol{\beta}'_i, \boldsymbol{\pi}'_i) \in \mathcal{B}_i$:

$$
\begin{aligned}
&\left| \pi_{ij} \nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) - \pi'_{ij} \nabla_{x_{ij}} u_i(\boldsymbol{\beta}'_i \oslash \boldsymbol{\pi}'_i) \right| \\
&= \left| \pi_{ij} \left(\nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) - \nabla_{x_{ij}} u_i(\boldsymbol{\beta}'_i \oslash \boldsymbol{\pi}'_i) \right) - \nabla_{x_{ij}} u_i(\boldsymbol{\beta}'_i \oslash \boldsymbol{\pi}'_i) \left(\pi_{ij} - \pi'_{ij} \right) \right| \\
&\leq \left| \pi_{ij} \left(\nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) - \nabla_{x_{ij}} u_i(\boldsymbol{\beta}'_i \oslash \boldsymbol{\pi}'_i) \right) \right| + \left| \nabla_{x_{ij}} u_i(\boldsymbol{\beta}'_i \oslash \boldsymbol{\pi}'_i) \left(\pi_{ij} - \pi'_{ij} \right) \right|
\end{aligned}
$$

$$
\begin{aligned}
&= \pi_{ij} \left| \nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) - \nabla_{x_{ij}} u_i(\boldsymbol{\beta}'_i \oslash \boldsymbol{\pi}'_i) \right| + \left| \nabla_{x_{ij}} u_i(\boldsymbol{\beta}'_i \oslash \boldsymbol{\pi}'_i) \right| \left| \pi_{ij} - \pi'_{ij} \right| \\
&\leq \pi_{ij} \left| \nabla_{x_{ij}} u_i(\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) - \nabla_{x_{ij}} u_i(\boldsymbol{\beta}'_i \oslash \boldsymbol{\pi}'_i) \right| + \ell \left| \pi_{ij} - \pi'_{ij} \right| \\
&\leq \pi_{ij} \lambda \left| (\boldsymbol{\beta}_i \oslash \boldsymbol{\pi}_i) - (\boldsymbol{\beta}'_i \oslash \boldsymbol{\pi}'_i) \right| - \ell \left| \pi_{ij} - \pi'_{ij} \right| \\
&\leq \pi_{ij} \lambda \left\| (\boldsymbol{\beta}_i - \boldsymbol{\beta}'_i) \oslash \boldsymbol{\pi}_i - \boldsymbol{\beta}'_i \odot \left(\mathbf{1}_m \oslash \boldsymbol{\pi}_i - \mathbf{1}_m \oslash \boldsymbol{\pi}'_i \right) \right\| + \ell \left| \pi_{ij} - \pi'_{ij} \right| \\
&\leq \pi_{ij} \lambda \left\| (\boldsymbol{\beta}_i - \boldsymbol{\beta}'_i) \oslash \boldsymbol{\pi}_i \right\| + \pi_{ij} \lambda \left\| \boldsymbol{\beta}'_i \odot \left(\mathbf{1}_m \oslash \boldsymbol{\pi}_i - \mathbf{1}_m \oslash \boldsymbol{\pi}'_i \right) \right\| + \ell \left| \pi_{ij} - \pi'_{ij} \right| \\
&\leq \frac{\pi_{ij} \lambda}{\min_{k \in [m]} \pi_{ik}} \left\| \boldsymbol{\beta}_i - \boldsymbol{\beta}'_i \right\| + \pi_{ij} \lambda \left\| \boldsymbol{\beta}'_i \odot \left(\mathbf{1}_m \oslash \boldsymbol{\pi}_i - \mathbf{1}_m \oslash \boldsymbol{\pi}'_i \right) \right\| + \ell \left| \pi_{ij} - \pi'_{ij} \right| \\
&\leq \frac{\pi_{ij} \lambda}{\min_{k \in [m]} \pi_{ik}} \left\| \boldsymbol{\beta}_i - \boldsymbol{\beta}'_i \right\| + \pi_{ij} \lambda \| \boldsymbol{\beta}'_i \| \left\| \mathbf{1}_m \oslash \boldsymbol{\pi}_i - \mathbf{1}_m \oslash \boldsymbol{\pi}'_i \right\| + \ell \left| \pi_{ij} - \pi'_{ij} \right| \\
&= \frac{\pi_{ij} \lambda}{\min_{k \in [m]} \pi_{ik}} \left\| \boldsymbol{\beta}_i - \boldsymbol{\beta}'_i \right\| + \pi_{ij} \lambda \| \boldsymbol{\beta}'_i \| \left\| \left(\boldsymbol{\pi}_i - \boldsymbol{\pi}'_i \right) \oslash \left(\boldsymbol{\pi}_i \odot \boldsymbol{\pi}'_i \right) \right\| + \ell \left| \pi_{ij} - \pi'_{ij} \right| \\
&\leq \frac{\pi_{ij} \lambda}{\min_{k \in [m]} \pi_{ik}} \left\| \boldsymbol{\beta}_i - \boldsymbol{\beta}'_i \right\| + \frac{\pi_{ij} \lambda}{\left\{ \min_{k \in [m]} \pi_{ik} \pi'_{ik} \right\}} \| \boldsymbol{\beta}'_i \| \left\| \boldsymbol{\pi}_i - \boldsymbol{\pi}'_i \right\| + \ell \left| \pi_{ij} - \pi'_{ij} \right| \\
&\leq \frac{\pi_{ij} \lambda}{\min_{k \in [m]} \pi_{ik}} \left\| \boldsymbol{\beta}_i - \boldsymbol{\beta}'_i \right\| + \frac{\pi_{ij} \lambda}{\left\{ \min_{k \in [m]} \pi_{ik} \pi'_{ik} \right\}} \| \boldsymbol{\beta}'_i \| \left\| \boldsymbol{\pi}_i - \boldsymbol{\pi}'_i \right\| + \ell \left\| \boldsymbol{\pi}_i - \boldsymbol{\pi}'_i \right\| \\
&\leq \frac{\pi_{\max} \lambda}{\pi_{\min}} \left\| \boldsymbol{\beta}_i - \boldsymbol{\beta}'_i \right\| + \frac{\pi_{\max} \lambda}{\pi_{\min}^2} \| \beta_{\max} \| \left\| \boldsymbol{\pi}_i - \boldsymbol{\pi}'_i \right\| + \ell \left\| \boldsymbol{\pi}_i - \boldsymbol{\pi}'_i \right\| \\
&\leq \max \left(\frac{\pi_{\max} \lambda}{\pi_{\min}}, \frac{\pi_{\max} \lambda}{\pi_{\min}^2}, \ell \right) \left\| (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) - (\boldsymbol{\beta}'_i, \boldsymbol{\pi}'_i) \right\|
\end{aligned}
$$

Putting it all together, for all consumers $i \in [n]$,

$$
\begin{aligned}
&\left\| \nabla u'_i(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) - \nabla u'_i(\boldsymbol{\beta}'_i, \boldsymbol{\pi}'_i) \right\| \\
&\leq \sum_{j \in [m]} \left\| \nabla_{\beta_{ij}} u'_i(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) - \nabla_{\beta_{ij}} u'_i(\boldsymbol{\beta}'_i, \boldsymbol{\pi}'_i) \right\| + \sum_{j \in [m]} \left\| \nabla_{\pi_{ij}} u'_i(\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) - \nabla_{\pi_{ij}} u'_i(\boldsymbol{\beta}'_i, \boldsymbol{\pi}'_i) \right\| \\
&\leq \sum_{j \in [m]} \max \left(\frac{\lambda}{\pi_{\min}^2}, \frac{\beta'_{\max} \lambda}{\pi_{\min}^2}, \frac{\ell}{\pi_{\min}^2} \right) \left\| (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) - (\boldsymbol{\beta}'_i, \boldsymbol{\pi}'_i) \right\| \\
&\quad + \sum_{j \in [m]} \max \left(\frac{\pi_{\max} \lambda}{\pi_{\min}}, \frac{\pi_{\max} \lambda}{\pi_{\min}^2}, \ell \right) \left\| (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) - (\boldsymbol{\beta}'_i, \boldsymbol{\pi}'_i) \right\| \\
&= m \max \left(\frac{\lambda}{\pi_{\min}^2}, \frac{\beta'_{\max} \lambda}{\pi_{\min}^2}, \frac{\ell}{\pi_{\min}^2} \right) \left\| (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) - (\boldsymbol{\beta}'_i, \boldsymbol{\pi}'_i) \right\| + m \max \left(\frac{\pi_{\max} \lambda}{\pi_{\min}}, \frac{\pi_{\max} \lambda}{\pi_{\min}^2}, \ell \right) \left\| (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) - (\boldsymbol{\beta}'_i, \boldsymbol{\pi}'_i) \right\| \\
&= m \left[\max \left(\frac{\lambda}{\pi_{\min}^2}, \frac{\beta'_{\max} \lambda}{\pi_{\min}^2}, \frac{\ell}{\pi_{\min}^2} \right) + \max \left(\frac{\pi_{\max} \lambda}{\pi_{\min}}, \frac{\pi_{\max} \lambda}{\pi_{\min}^2}, \ell \right) \right] \left\| (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) - (\boldsymbol{\beta}'_i, \boldsymbol{\pi}'_i) \right\|
\end{aligned}
$$

Remark 10.3.4 [Boundedness of bid space]**.**

Recall that in the trading post pseudo-game, the action spaces $\mathcal{B}_i$ of the players $i \in [n]$, are compact by definition. As such, $\pi_{\min}, \pi_{\max}, \beta_{\max}$ are all guaranteed to exist.

With sufficient conditions that ensure the Lipschitz-smoothness of the trading post pseudo-game, we combine Lemma 10.3.2, Lemma 10.3.5, Lemma 10.3.6, and Theorem 4.3.1 to arrive at the following theorem.

Theorem 10.3.1 [Convergence of Mirror Extratrade Learning Dynamics]**.**

Consider a concave pure exchange economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, where for all players $i \in [n]$, u_i is ℓ-Lipschitz-continuous and λ-Lipschitz-smooth. Let $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$ be the trading post pseudo-game (Definition 10.2.1) associated with the pure exchange economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, and h a 1-strongly-convex and κ-Lipschitz-smooth kernel function. Define:

$$\pi_{\min} \doteq \min\{\pi_{ij} : (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \in \mathcal{B}_i, \forall i \in [n], j \in [m]\} \enspace ,$$

$$\pi_{\max} \doteq \max\{\pi_{ij} : (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \in \mathcal{B}_i, \forall i \in [n], j \in [m]\} \enspace ,$$

$$\beta_{\max} \doteq \max\{\beta_{ij} : (\boldsymbol{\beta}_i, \boldsymbol{\pi}_i) \in \mathcal{B}_i, \forall i \in [n], j \in [m]\} \enspace ,$$

and

$$\lambda' \doteq m \left[\max\left(\frac{\lambda}{\pi_{\min}^2}, \frac{\beta'_{\max}\lambda}{\pi_{\min}^2}, \frac{\ell}{\pi_{\min}^2} \right) + \max\left(\frac{\pi_{\max}\lambda}{\pi_{\min}}, \frac{\pi_{\max}\lambda}{\pi_{\min}^2}, \ell \right) \right] \enspace .$$

Consider the mirror extratrade dynamics, i.e., the mirror extragradient learning dynamics (Algorithm 7) applied to the trading post pseudo-game, run on the trading post pseudo-game $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$ with the kernel function h, a step size $\eta \in \left(0, \frac{1}{\sqrt{2}\lambda'}\right]$, and any time horizon $\tau \in \mathbb{N}$. The output sequence $\{(\boldsymbol{\beta}^{(t+0.5)}, \boldsymbol{\pi}^{(t+0.5)}), (\boldsymbol{\beta}^{(t)}, \boldsymbol{\pi}^{(t)})\}_t$ satisfies the following: If $(\boldsymbol{\beta}^{(\tau)}_{\text{best}}, \boldsymbol{\pi}^{(\tau)}_{\text{best}}) \in \arg\min_{(\boldsymbol{\beta}^{(k)}, \boldsymbol{\pi}^{(k)}): k=0,\dots,\tau} \mathrm{div}_h((\boldsymbol{\beta}^{(k+0.5)}, \boldsymbol{\pi}^{(k+0.5)}), (\boldsymbol{\beta}^{(k)}, \boldsymbol{\pi}^{(k)}))$, then for of $\tau \in O(1/\varepsilon^2)$, $(\boldsymbol{\beta}^{(\tau)}_{\text{best}}, \boldsymbol{\pi}^{(\tau)}_{\text{best}})$ is a ε-first-order VE of $(\mathcal{A}, \boldsymbol{g}, \boldsymbol{u}')$. In addition, if we define $\boldsymbol{p}^{(t)} \doteq \sum_{k \in [n]} \boldsymbol{\beta}^{(t)}_k$ and $\boldsymbol{x}^{(t)}_i \doteq \boldsymbol{\beta}^{(t)}_i \oslash \boldsymbol{\pi}^{(t)}_i$, then $\lim_{t \to \infty} (\boldsymbol{x}^{(t)}, \boldsymbol{p}^{(t)}) = (\boldsymbol{x}^*, \boldsymbol{p}^*)$ is an Arrow-Debreu equilibrium of $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$.

10.4 Merit Function Methods for Arrow-Debreu Economies

10.4.1 Merit Functions for Arrow-Debreu Economies

In this section, we investigate the computation of Arrow-Debreu equilibrium beyond pure exchange economies. As the efficient computation of an Arrow-Debreu equilibrium seems out of reach beyond pure exchange economies, we will loosen our aim of computing an Arrow-Debreu equilibrium to the computation of prices and consumptions that satisfy necessary conditions to be an Arrow-Debreu equilibrium. Our approach to computing such prices and consumptions will be to introduce two separate polynomial-time first- and second-order methods, which we will derive via merit function minimization.

In particular, we will consider the Arrow-Debreu pseudo-game whose set of GNE is equal to the set of Arrow-Debreu competitive equilibria of the associated Arrow-Debreu economy, and then apply the merit function methods for pseudo-games we derived in Section 9.4 to compute an action profile that satisfies necessary condition to be a GNE of the Arrow-Debreu pseudo-game, thus also resulting in prices and consumptions that satisfy necessary conditions to be an Arrow-Debreu equilibrium.

First, let's recall the Arrow-Debreu pseudo-game (Definition 10.2.1), which consists of the following $n+1$ simultaneous optimization problems:

$$\forall i \in [n], \qquad \max_{\boldsymbol{x}_i \in \mathcal{X}_i' : \boldsymbol{x}_i \cdot \boldsymbol{p} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}} u_i(\boldsymbol{x}_i) \qquad \Bigg| \qquad \max_{\boldsymbol{p} \in \Delta_m} \boldsymbol{p} \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i - \sum_{i \in [n]} \boldsymbol{e}_i \right)$$

Now, as this pseudo-game does not have jointly convex constraints, a VE is not guaranteed to exist; as such, none of the methods we derived in Section 9.4 are applicable. To overcome this difficulty, we instead leverage Theorem 9.2.3, which allows us to convert the Arrow-Debreu pseudo-game into a $2n+1$ player *game*, and apply our merit function methods to solve this game.

Definition 10.4.1 [Arrow-Debreu Game]**.**

Given an Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, we define the associated $(2n+1)$-player **Arrow-**

Debreu game $(n+1, 1, \mathcal{A}, \boldsymbol{u}')$, denoted $(\mathcal{A}, \boldsymbol{u}')$ when n is clear from context, in which the first n players are called "consumers", the players n through $2n$ are called the "shadow consumers," and the $(2n+1)^{th}$ player is called the "auctioneer," and where:

(Action spaces) For all consumers $i \in [n]$, $\mathcal{A}_i \doteq \mathcal{X}'_i \doteq \left\{\boldsymbol{x}_i \mid \sum_{k \in [n]} \boldsymbol{x}_k \leq \sum_{k \in [n]} \boldsymbol{e}_k, \forall \boldsymbol{x} \in \mathcal{X}\right\}$,
for all shadow consumers $i \in [n, 2n]$, $\mathcal{A}_i \doteq \Lambda_i \subseteq \mathbb{R}_+$,
and for the auctioneer $\mathcal{A}_{2n+1} \doteq \Delta_m$

(Payoffs) For all consumers $i \in [n]$, $u'_i(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{p}) \doteq u_i(\boldsymbol{x}_i) + \lambda_i (\boldsymbol{e}_i \cdot \boldsymbol{p} - \boldsymbol{x}_i \cdot \boldsymbol{p})$,
for all shadow consumers $i \in [n, 2n]$, $u'_i(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{p}) \doteq -u_i(\boldsymbol{x}_i) - \lambda_i (\boldsymbol{e}_i \cdot \boldsymbol{p} - \boldsymbol{x}_i \cdot \boldsymbol{p})$,
and for the auctioneer, $u'_{n+1}(\boldsymbol{x}, \boldsymbol{\lambda}, \boldsymbol{p}) \doteq \boldsymbol{p} \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i - \sum_{i \in [n]} \boldsymbol{e}_i\right)$.

The Arrow-Debreu game can succinctly be represented by the following $2n+1$ simultaneous optimization problems:

$$\forall i \in [n], \qquad \max_{\boldsymbol{x}_i \in \mathcal{X}'_i} \min_{\lambda_i \geq 0} u_i(\boldsymbol{x}_i) + \lambda_i (\boldsymbol{e}_i \cdot \boldsymbol{p} - \boldsymbol{x}_i \cdot \boldsymbol{p}) \qquad \Big| \max_{\boldsymbol{p} \in \Delta_m} \boldsymbol{p} \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i - \sum_{i \in [n]} \boldsymbol{e}_i\right)$$

We note that a naive choice for the action spaces of the shadow consumers is for all $i \in [n, 2n]$, $\Lambda_i \doteq \mathbb{R}_+$. To ensure polynomial-time convergence of our algorithm, however, it will be necessary to instead choose a non-empty and compact Λ_i, to ensure that the players' utility functions are Lipschitz-smooth. While we will not delve in the details of how to choose such action spaces—to keep our exposition simple, we will only present informal expositions of our theorems—we note that under Slater's condition, which is satisfied in quasiconcave Arrow-Debreu economies, it is possible to define such a set. For additional details, we refer the reader to Nedić and Ozdaglar (2009) and Nedic and Ozdaglar (2009), as well as Definition 9.2.6.

Now, notice that in quasiconcave Arrow-Debreu economies, Slater's condition is guaranteed to hold. As such, we have the following corollary of Theorem 9.2.3 and Lemma 10.2.1:

Corollary 10.4.1.

Given a quasiconcave Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, consider the associated Arrow-Debreu game $(\mathcal{A}, \boldsymbol{u}')$. For any Arrow-Debreu equilibrium $(\boldsymbol{x}^*, \boldsymbol{p}^*)$ of the Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, there exists $\boldsymbol{\lambda}^* \in \Lambda$ s.t. $(\boldsymbol{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{p}^*)$ is a Nash equilibrium of the Arrow-Debreu game $(\mathcal{A}, \boldsymbol{u}')$.

Conversely, the consumptions and prices $(\boldsymbol{x}^*, \boldsymbol{p}^*)$ of any Nash equilibrium $(\boldsymbol{x}^*, \boldsymbol{\lambda}^*, \boldsymbol{p}^*)$ of the Arrow-Debreu game $(\mathcal{A}, \boldsymbol{u}')$ correspond to an Arrow-Debreu equilibrium of the Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$.

10.4.2 First-Order Market Dynamics for Merit Function Minimization

With this corollary in hand, we can now apply REDA (Algorithm 9) to compute a stationary point of the regularized exploitability associated with the Arrow-Debreu game, or alternatively Algorithm 10 to compute a stationary point of the variational exploitability associated with the Arrow-Debreu game. Our first theorem is a corollary of Theorem 9.4.3, which we present informally to avoid burdening our exposition with highly involved bounds which depend on the Lipschitz-smoothness constant of the utility functions of the consumers and the radius of the action space of the shadow consumers.

Theorem 10.4.1 [Convergence of REDA in the Arrow-Debreu Game]**.**
Given an Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, consider the associated Arrow-Debreu game $(\mathcal{A}, \boldsymbol{u}')$, and assume that for all consumers $i \in [n]$, u_i is Lipschitz-smooth, and for all shadow consumers $i \in [n, 2n]$, the action spaces $\mathcal{A}_i \doteq \Lambda_i$ are non-empty, compact, convex and contain the Nash equilibrium actions of the shadow consumers. For $\varepsilon, \alpha > 0$, if φ_α be the α-regularized exploitability associated with the Arrow-Debreu game, h is a 1-strongly-convex kernel function, and $\boldsymbol{x}^{(0)}, \boldsymbol{\lambda}^{(0)}, \boldsymbol{p}^{(0)} \in \mathcal{A}$ are some initial actions, then for some appropriate choice of $\eta > 0$ and $\tau \in O(1/\varepsilon^2)$, the regularized extragradient descent ascent algorithm (REDA; Algorithm 9) is guaranteed to compute an ε-stationary point of the α-regularized exploitability φ_α.

10.4.3 Second-Order Market Dynamics for Merit Function Minimization

Similar to Theorem 10.4.1, we also obtain the following corollary of Theorem 9.6.2 for the convergence of the mirror variational learning dynamics when applied to the Arrow-Debreu game.

Theorem 10.4.2.

Given an Arrow-Debreu economy $(\mathcal{X}, \boldsymbol{e}, \boldsymbol{u})$, consider the associated Arrow-Debreu game $(\mathcal{A}, \boldsymbol{u}')$, and assume that for all consumers $i \in [n]$, u_i is Lipschitz-smooth, and for all shadow consumers $i \in [n, 2n]$, the action spaces $\mathcal{A}_i \doteq \lambda_i$ are non-empty, compact, convex and contain the Nash equilibrium actions of the shadow consumers. For $\varepsilon, \alpha > 0$, if Ξ_α is the α-variational exploitability associated with the Arrow-Debreu game, h is a 1-strongly-convex kernel function, and $\boldsymbol{x}^{(0)}, \boldsymbol{\lambda}^{(0)}, \boldsymbol{p}^{(0)} \in \mathcal{A}$ are some initial actions, then for some appropriate choice of $\eta > 0$ and $\tau \in O(1/\varepsilon)$, the mirror variational learning dynamics is guaranteed to compute an ε-stationary point of the α-variational exploitability Ξ_α.

Part III

Markov Pseudo-Games and Radner Economies

Chapter 11

Scope and Motivation

11.1 Scope

Part III of this thesis is divided into two chapters. In Chapter 12, we[1] introduce our model of Markov pseudo-games (Section 12.1) and define appropriate solution concepts, i.e., equilibria, in Section 12.2, where we establish their existence under suitable conditions. We then present a gradient descent-ascent-based reinforcement learning algorithm (Algorithm 11; TTSSGDA), which provably converges to a solution satisfying the necessary equilibrium conditions via a coupled min-max optimization formulation of the problem (Section 12.3).

In Chapter 13, we apply our theory to Radner economies (or infinite-horizon Markov exchange economies). First, we formulate static exchange economies, i.e., spot markets, as generalized (one-shot) games. Next, we develop infinite-horizon Radner economies by modeling them as instances of our Markov pseudo-games framework. We then invoke our main theorems for Markov pseudo-games to establish the existence of recursive Radner equilibria in Radner economies—the first such result to our knowledge—as well as the convergence of TTSSGDA to this equilibrium. Finally, we present experiments confirming the accuracy and efficiency of TTSSGDA in practice.

[1]The work in Part III was developed in collaboration with Sadie Zhao, Yiling Chen, and Amy Greenwald.

11.2 Motivation

In 1896, Léon Walras formulated a mathematical model of markets as a system for resource allocation comprising supply and demand functions that map values for resources, called **prices**, to quantities of resources—*ceteris paribus*, i.e., all else being equal. Walras argued that any market would eventually settle into a steady state, which he called **competitive** (nowadays, also called **Walrasian**) **equilibrium**, as a collection of prices and associated supply and demand such that the demand is **feasible**, i.e., the demand for each resource is less than or equal to its supply, and **Walras' law** holds, i.e., the value of the supply is equal to the value of the demand. Unlike in Walras' model, real-world markets do not exist in isolation but are part of an **economy**. Indeed, the supply and demand of resources in one market depend not only on prices in that market, but also on the supply and demand of resources in other markets. If every market in an economy is simultaneously at a competitive equilibrium, Walras' law holds for the economy as a whole; this steady state, now a property of the economy, is called a **general equilibrium**.

Beyond Walras' early forays into competitive equilibrium analysis, foremost to the development of the theory of general equilibrium was the introduction of a broad mathematical framework for modeling economies, which is known today as the **Arrow-Debreu competitive economy** (Arrow and Debreu, 1954). In this same paper, Arrow and Debreu developed their seminal game-theoretic model, namely (quasi)concave pseudo-games, and proved the existence of generalized Nash equilibrium in this model. Since this game-theoretic model is sufficiently rich to capture Arrow-Debreu economies, they obtained as a corollary the existence of general equilibrium in these economies.

In their model, Arrow and Debreu posit a set of resources, modeled as commodities, each of which is assigned a price; a set of consumers, each choosing a quantity of each commodity to consume in exchange for their endowment; and a set of firms, each choosing a quantity of each commodity to produce, with prices determining **(aggregate) demand**, i.e., the sum of the consumptions across all consumers, and **(aggregate) supply**, i.e., the sum of endowments and productions across all consumers and firms, respectively. This model is **static**, as it comprises only a single period model, but it is nonetheless rich, as commodities can be state and time contingent, with each one representing a good or service which can be bought or sold in a single time period, but that encodes delivery opportunities at a finite number of distinct points in time. Following Arrow and Debreu's seminal existence result, the literature slowly turned away from static economies, such as Arrow-Debreu competitive economies, which do not *explicitly* involve time and uncertainty.

Arrow and Debreu's model fails to provide a comprehensive account of the economic activity observed in the real world, especially that which is designed to account for *time* and *uncertainty*. Chief among these activities are **asset markets**, which allow consumers and firms to insure themselves against uncertainty about future states of the world. Indeed, while static economies with state- and time-contingent commodities can *implicitly* incorporate time and uncertainty, the assumption that a complete set of state- and time-contingent commodities are available at the time of trade is highly unrealistic. Arrow (1964) thus proposed to enhance the Arrow-Debreu competitive economy with **assets** (or **securities** or **stocks**),[2] i.e., contracts between two consumers, which promise the delivery of commodities by its seller to its buyer at a future date. In particular, Arrow introduced an asset type nowadays known as the ***numéraire* Arrow security**, which transfers one unit of a designated commodity used as a unit of account—**the numéraire**—when a particular state of the world is observed, and nothing otherwise. As the numéraire is often interpreted as money, assets which deliver only some amount of the numéraire, are called **financial assets**.

Formally, Arrow considered a **two-step stochastic exchange economy**. In the initial state, consumers can buy or sell *numéraire* Arrow securities in a **financial asset market**. Following these trades, the economy stochastically transitions to one of finitely many other states in

[2]Some authors (e.g., Geanakoplos (1990)) distinguish between assets, stocks, and securities, instead defining securities (respectively, stocks) as those assets which are defined exogenously (respectively, endogenously), e.g., government bonds (respectively, company stocks). As this distinction makes no mathematical difference to our results, and is only relevant to stylized models, we make no such distinction.

which consumers receive returns on their initial investment and participate in a **spot market**, i.e., a market for the immediate delivery of commodities, modeled as a static exchange economy—which, for our purposes, is better called an **exchange market**.[3] A general equilibrium of this economy is then simply prices for financial assets *and* commodities, which lead to a feasible allocation of all resources (i.e., financial assets and spot market commodities) that satisfies Walras' law.

Arrow (1964) demonstrated that the general equilibrium consumptions of an exchange economy with state- and time-contingent commodities can be implemented by the general equilibrium spot market consumptions of a two-step stochastic economy with a considerably smaller, yet **complete** set of *numéraire* Arrow securities, i.e., a set of securities available for purchase in the first period that allow consumers to transfer wealth to *all* possible states of the world that can be realized in the second period. In conjunction with the welfare theorems (Debreu, 1951a; Arrow, 1951b), this result implies that economies with **complete financial asset markets**, i.e., economies with such a complete set of securities, achieve a Pareto-optimal allocation of commodities by ensuring optimal risk-bearing via financial asset markets; and conversely, any Pareto-optimal allocation of commodities in economies with time and uncertainty can be realized as a competitive equilibrium of a complete financial asset market.

Arrow's contributions led to the development of a new class of general equilibrium models, namely **stochastic economies** (or **dynamic stochastic general equilibrium—DSGE—models**) (Geanakoplos, 1990).[4] At a high-level, these models comprise a sequence of world states and spot markets, which are linked across time by asset markets, with each next state of the world (respectively, spot market) determined by a stochastic process that is independent of market interactions (respectively, dependent only on their asset purchase)

[3]An (Arrow-Debreu) exchange economy is simply an (Arrow-Debreu) competitive economy without firms. Historically, for simplicity, it has become standard practice *not* to model firms, as most, if not all, results extend directly to settings with firms. In line with this practice, we do not model firms, but note that our results and methods also extend directly to settings that include firms.

[4]As these models incorporate both time and uncertainty, they are often referred to as dynamic stochastic general equilibrium models. Nonetheless, we opt for the stochastic economy nomenclature, because, as we demonstrate in this paper, these economies can be seen as instances of (generalized) stochastic games.

in the current state. Mathematically, the key difference between a static and a stochastic economy is that consumers in a stochastic economy face a collection of budget constraints, one per time-step, rather than only one. Indeed, Arrow (1964)'s proof that general equilibrium consumptions in stochastic complete economies are equivalent to general equilibrium consumptions in static state- and time-contingent commodity economies relies on proving that the many budget constraints in a complete stochastic economy can be reduced to a single one.

Stochastic economies were introduced to model arbitrary finite time horizons (Radner, 1968) and a variety of risky asset classes (e.g., stocks (Diamond, 1967), risky assets (Lintner, 1975), derivatives (Black and Scholes, 1973), capital assets (Mossin, 1966), debts (Modigliani and Miller, 1958) etc.), eventually leading to the emergence of **stochastic economies with incomplete asset markets** (Magill and Shafer, 1991; Magill and Quinzii, 2002; Geanakoplos, 1990), or colloquially, **(incomplete) stochastic economies**.[5] Unlike in Arrow's stochastic economy, the asset market is not complete in such economies, so consumers cannot necessarily insure themselves against all future world states.

The archetypal stochastic economy is the **Radner stochastic exchange economy**, deriving its name from Radner's proof of existence of a general equilibrium in his model (Radner, 1972). Radner's economy is a finite-horizon stochastic economy comprising a sequence of spot markets, linked across time by asset markets. At each time period, a finite set of consumers observe a world state and trade in an asset market and a spot market, modeled as an exchange market. Each **asset market** comprises **assets**, modelled as time-contingent **generalized Arrow securities**, which specify quantities of the commodities the seller is obliged to transfer to its buyer, should the relevant state of the economy be realized at some specified future time.[6] Consumers can buy and sell assets, thereby transferring their

[5]While many authors have called these models incomplete economies (Geanakoplos, 1990; Magill and Quinzii, 2002; Magill and Shafer, 1991), these models capture both incomplete and complete asset markets. In contrast, we refer to stochastic economies with incomplete or complete asset markets as **stochastic economies**, adding the (in)complete epithet only when necessary to indicate that the asset market is (in)complete.

[6]Here, Arrow securities are "generalized" in the sense that they can deliver different quantities of *many* commodities at different states of the world, rather than only one unit of a commodity at only one state of the

wealth across time, all the while insuring themselves against uncertainty about the future. The canonical solution concept for stochastic economies, **Radner equilibrium** (also called **sequential competitive equilibrium**[7] (Mas-Colell et al., 1995), **rational expectations equilibrium** (Radner, 1979), and **general equilibrium with incomplete markets** (Geanakoplos, 1990)), is a collection of history-dependent prices for commodities and assets, as well as history-dependent consumptions of commodities and portfolios of assets, such that, for all histories, the aggregate demand for commodities and the **aggregate demand for assets** (i.e., the total number of assets bought) are feasible and satisfy Walras' law.

In spite of substantial interest in stochastic economies among microeconomists throughout the 1970s, the literature eventually trailed off, perhaps due to a seemingthe difficulty inof proving existence of a general equilibrium in simple economies with incomplete asset markets that allow assets to be sold short (Geanakoplos, 1990), or to the lack of a second welfare theorem (Dreze, 1974; Hart, 1975). Financial and macroeconomists stepped up, however, with financial economists seeking to further develop the theoretical aspects of stochastic economies (see, for instance, Magill and Quinzii (2002)), and macroeconomists seeking practical methods by which to solve stochastic economies in order to determine the impact of various policy choices (via simulation; see, for instance, Sargent and Ljungqvist (2000)).

Radner economies are one of the new and interesting directions in this more recent work on stochastic economies. Infinite horizon models come with one significant difficulty that has no counterpart in a finite horizon model, namely the possibility for agents to run a **Ponzi scheme** via asset markets, in which they borrow but then indefinitely postpone repaying their debts by refinancing them continually, from one period to the next. From this perspective infinite horizon models represent very interesting objects of study, not only theoretically; it has also been argued that they are a better modeling paradigm

world. Although Arrow (1964) considered only *numéraire* securities, his theory was subsequently generalized to generalized Arrow securities (Geanakoplos, 1990).

[7] This terminology does not contradict the economy being at a competitive equilibrium, but rather indicates that at all times, the spot and asset markets are at a competitive equilibrium, hence implying the overall economy is at a general equilibrium.

for macroeconomists who employ simulations (Magill and Quinzii, 1994), because they facilitate the modeling of complex phenomena, such as asset bubbles (Huang and Werner, 2000), which can be impacted by economic policy decisions.

Magill and Quinzii (1994) introduced an extension of Radner's model to an infinite horizon setting, albeit with financial assets, and presented suitable assumptions under which a sequential competitive equilibrium is guaranteed to exist in this model. Progress on the computational aspects of stochastic economies has been slow, however, and mostly confined to finite horizon settings (see, Sargent and Ljungqvist (2000) and Volume 2 of Taylor and Woodford (1999) for a standard survey, and Fernández-Villaverde (2023) for a more recent entry-level survey of computational methods used by macroeconomists). Indeed, demands for novel computational methods for solving macroeconomic models, and theoretical frameworks in which to understand their computational complexity, have been repeatedly shared by macroeconomists (Taylor and Woodford, 1999). This gap in the literature points to a novel research opportunity; however, it is challenging for non-macroeconomists to approach these problems with their computational tools.

11.3 Contributions

In Chapter 12, we introduce Markov pseudo-games, and we establish the existence of (pure) **generalized Markov perfect equilibria (GMPE)** in concave Markov pseudo-games (Theorem 12.2.1). This result can be seen as a stochastic generalization of Arrow and Debreu (1954)'s existence result for (pure) generalized Nash equilibrium in concave pseudo-games (Facchinei and Kanzow, 2010a). It also implies the existence of pure (or deterministic) Markov perfect equilibria in a large class of continuous-action Markov games for which, to the best of our knowledge, existence was heretofore known only in mixed (or randomized) policies (Fink, 1964; Takahashi, 1964).

Although the computation of GMPE is PPAD-hard in general, because GMPE generalize Nash equilibrium, we reduce this computational problem to generative adversarial learning between a generator, who produces a candidate equilibrium policy profile, and an adversary, who produces a policy profile of best responses to the candidate equilibrium (Goktas et al., 2023a) (Observation 12.3.1). Assuming parameterized policies, and taking advantage of the recent progress on solving generative adversarial learning problems (e.g., (Lin et al., 2020; Daskalakis et al., 2020a)), we show that a policy profile that is a stationary point of the exploitability (i.e., the players' cumulative maximum regret) can be computed in polynomial time under mild assumptions (Theorem 12.3.1). This result implies that a policy profile that satisfies necessary first-order stationarity conditions for a GMPE in Markov pseudo-games with a bounded best-response mismatch coefficient (Lemma 12.3.4)—i.e., those Markov pseudo-games in which states explored by any GMPE are easily explored under the initial state distribution—can be computed in polynomial time, a result which is analogous known computational results for zero-sum Markov games (Daskalakis et al., 2020a). As our theoretical computational guarantees apply to policies represented by neural networks, we obtain the first, to our knowledge, deep reinforcement learning algorithm with theoretical guarantees for general-sum games.

In Chapter 13, we introduce an extension of Magill and Quinzii (1994)'s infinite horizon exchange economy, which we call the **Radner economy**. On the one hand, our model generalizes Magill and Quinzii's to a setting with arbitrary, not just financial, assets; on the other hand, we restrict the transition model to be Markov. The Markov restriction allows us to prove the existence of a **recursive Radner equilibrium (RRE)** (Mehra and Prescott, 1977) (Theorem 13.1.1). Our proof reformulates the set of RRE of any Radner economy as the set of GMPE of an associated generalized Markov game (Theorem 13.1.1). To our knowledge, ours is the first result of its kind for such a general setting, as previous recursive competitive equilibrium existence proofs were restricted to economies with one consumer (also called the representative agent), one commodity, or one asset (Mehra and Prescott, 1977; Prescott and Mehra, 1980). The aforementioned results allow us to conclude that a stationary point of the exploitability of the Markov pseudo-game associated with any Radner economy can be computed in polynomial time (Theorem 13.1.2).

Finally, in Section 13.2 we implement our policy gradient method in the form of a generative adversarial policy network (GAPNet), and use it to try to find RRE in three Radner economies with three different types of utility functions. Experimentally, we find that our GAPNet produces approximate equilibrium policies that are closer to GMPE than those produced by a standard macroeconomic baseline for solving stochastic economies.

Chapter 12

Markov Pseudo-Games

12.1 Background

12.1.1 Mathematical Background

Throughout, we adopt the following notational shorthand: If a function $f : \mathcal{X} \to \mathbb{R}$ is Lipschitz-continuous (resp. Lipschitz-smooth), for simplicity, we will denote its Lipschitz-continuity (resp. Lipschitz-smoothness) constant by $\ell_f \geq 0$ (resp. $\ell_{\nabla f} \geq 0$).

We will also require notions of stochastic convexity related to stochastic dominance of probability measures (Atakan, 2003b).

Definition 12.1.1 [Stochastic Convexity/Concavity]**.**
Given non-empty and convex parameter and outcome spaces $\mathcal{W}$ and $\mathcal{O}$ respectively, a conditional probability distribution $\boldsymbol{w} \mapsto \rho(\cdot \mid \boldsymbol{w}) \in \Delta(\mathcal{O})$ is said to be **stochastically convex** (resp. **stochastically concave**) in $\boldsymbol{w} \in \mathcal{W}$ if for all continuous, bounded, and convex (resp. concave) functions $v : \mathcal{O} \to \mathbb{R}$, $\lambda \in (0, 1)$, and $\boldsymbol{w}', \boldsymbol{w}^\dagger \in \mathcal{W}$ s.t. $\overline{\boldsymbol{w}} = \lambda \boldsymbol{w}' + (1-\lambda)\boldsymbol{w}^\dagger$, it holds that $\mathbb{E}_{O \sim \rho(\cdot \mid \overline{\boldsymbol{w}})} [v(O)] \leq$ (resp. $\geq$) $\lambda \, \mathbb{E}_{O \sim \rho(\cdot \mid \boldsymbol{w}')} [v(O)] + (1-\lambda) \, \mathbb{E}_{O \sim \rho(\cdot \mid \boldsymbol{w}^\dagger)} [v(O)]$.

12.1.2 Markov Pseudo-Games

We begin by developing our formal game model. The games we study are stochastic, in the sense of Shapley (1953), Fink (1964), and Takahashi (1964). Further, they are pseudo-games, in the sense of Arrow and Debreu (1954). Arrow and Debreu introduced pseudo-games to establish the existence of competitive equilibrium in their seminal model of an exchange economy, where an auctioneer sets prices that determine the consumers' budget sets, and hence their feasible consumptions. It is this dependency among the players' feasible actions that characterizes pseudo-games. We model stochastic pseudo-games, and dub them **Markov pseudo-games**, as the games are Markov in that the stochastic transitions depend only on the most recent state and player actions.

An **(infinite horizon discounted) Markov pseudo-game** $\mathcal{M} \doteq (n, m, \mathcal{S}, \mathcal{A}, \mathcal{X}, \boldsymbol{r}, \rho, \gamma, \mu)$ is an n-player dynamic game played over an infinite discrete time horizon. The game starts at time $t = 0$ in some initial state $S^{(0)} \sim \mu \in \Delta(\mathcal{S})$ drawn randomly from a set of states $\mathcal{S} \subseteq \mathbb{R}^l$. At this and each subsequent time period $t = 1, 2, \ldots$, the players encounter a state $\boldsymbol{s}^{(t)} \in \mathcal{S}$, in which each $i \in [n]$ simultaneously takes an **action** $\boldsymbol{a}_i^{(t)} \in \mathcal{X}_i(\boldsymbol{s}^{(t)}, \boldsymbol{a}_{-i}^{(t)})$ from a **set of feasible actions** $\mathcal{X}_i(\boldsymbol{s}^{(t)}, \boldsymbol{a}_{-i}^{(t)}) \subseteq \mathcal{A}_i \subseteq \mathbb{R}^m$, determined by a **feasible action correspondence** $\mathcal{X}_i : \mathcal{S} \times \mathcal{A}_{-i} \rightrightarrows \mathcal{A}_i$, which takes as input the current state $\boldsymbol{s}^{(t)}$ and the other players' actions $\boldsymbol{a}_{-i}^{(t)} \in \mathcal{A}_{-i}$, and outputs a subset of the ith player's action space $\mathcal{A}_i$. We define $\mathcal{X}(\boldsymbol{s}, \boldsymbol{a}) \doteq \bigtimes_{i \in [n]} \mathcal{X}_i(\boldsymbol{s}, \boldsymbol{a}_{-i})$.

Once the players have taken their actions $\boldsymbol{a}^{(t)} \doteq (\boldsymbol{a}_1^{(t)}, \ldots, \boldsymbol{a}_n^{(t)})$, each player $i \in [n]$ receives a **reward** $r_i(\boldsymbol{s}^{(t)}, \boldsymbol{a}^{(t)})$ given by a **reward function** $\boldsymbol{r} : \mathcal{S} \times \mathcal{A} \to \mathbb{R}^n$, after which the game either ends with probability $1-\gamma$, where $\gamma \in (0, 1)$ is called the **discount factor**,[1] or continues on to time period $t + 1$, transitioning to a new state $S^{(t+1)} \sim \rho(\cdot \mid \boldsymbol{s}^{(t)}, \boldsymbol{a}^{(t)})$, according to a **transition** probability function $\rho : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \to [0, 1]$, where $\rho(\boldsymbol{s}^{(t+1)} \mid \boldsymbol{s}^{(t)}, \boldsymbol{a}^{(t)}) \in [0, 1]$ denotes the probability of transitioning to state $\boldsymbol{s}^{(t+1)} \in \mathcal{S}$ from state $\boldsymbol{s}^{(t)} \in \mathcal{S}$ when action profile $\boldsymbol{a}^{(t)} \in \mathcal{A}$ is played.

Our focus is on continuous-state and continuous-action Markov pseudo-games, where the state and action spaces are non-empty and compact, and the reward functions are continuous and bounded in each of $\boldsymbol{s}$ and $\boldsymbol{a}$, holding the other fixed.

[1] Our results generalize to settings with per-player discount factors $\gamma_i \in (0, 1)$, where the discount rates express the players' intertemporal preferences over game outcomes at each time-step. .

A **history** $\boldsymbol{h} \in \mathcal{H}^\tau \doteq (\mathcal{S} \times \mathcal{A})^\tau \times \mathcal{S}$ of length $\tau \in \mathbb{N}$ is a sequence of states and action profiles $\boldsymbol{h} = ((\boldsymbol{s}^{(t)}, \boldsymbol{a}^{(t)})_{t=0}^{\tau-1}, \boldsymbol{s}^{(\tau)})$ s.t. a history of length 0 corresponds only to the initial state of the game. For any history $\boldsymbol{h} = ((\boldsymbol{s}^{(t)}, \boldsymbol{a}^{(t)})_{t=0}^{\tau-1}, \boldsymbol{s}^{(\tau)})$ of length $\tau \in \mathbb{N}$, we denote by $\boldsymbol{h}_{:\tau'}$ the first $\tau' \in [0 : \tau]$ steps of $\boldsymbol{h}$, i.e., $\boldsymbol{h}_{:\tau'} = ((\boldsymbol{s}^{(t)}, \boldsymbol{a}^{(t)})_{t=0}^{\tau'-1}, \boldsymbol{s}^{(\tau')})$.

Overloading notation, we define the **history space** $\mathcal{H} \doteq \bigcup_{\tau=0}^{\infty} \mathcal{H}^\tau$. For any player $i \in [n]$, a **policy** $\boldsymbol{\pi}_i : \mathcal{H} \to \mathcal{A}_i$ is a mapping from histories of any length to i's space of (pure) actions. We define the space of all (deterministic) policies as $\mathcal{P}_i \doteq \{\boldsymbol{\pi}_i : \mathcal{H} \to \mathcal{A}_i\}$.[2] A **Markov policy** (Maskin and Tirole, 2001) $\boldsymbol{\pi}_i$ is a policy s.t. $\boldsymbol{\pi}_i(\boldsymbol{s}^{(\tau)}) = \boldsymbol{\pi}_i(\boldsymbol{h}_{:\tau})$, for all histories $\boldsymbol{h} \in \mathcal{H}^\tau$ of length $\tau \in \mathbb{N}_+$, where $\boldsymbol{s}^{(\tau)}$ denotes the final state of history $\boldsymbol{h}$. As Markov policies are only state-contingent, we can compactly represent the space of all Markov policies for player $i \in [n]$ as $\mathcal{P}_i^{\text{markov}} \doteq \{\boldsymbol{\pi}_i : \mathcal{S} \to \mathcal{A}_i\}$.

Fixing player $i \in [n]$ and $\boldsymbol{\pi}_{-i} \in \mathcal{P}_{-i}$, given history $\boldsymbol{h} \in \mathcal{H}^\tau$, we define the **feasible policy correspondence**

$$\mathcal{F}_i(\boldsymbol{\pi}_{-i}) \doteq \{\boldsymbol{\pi}_i \in \mathcal{P}_i \mid \forall \boldsymbol{h} \in \mathcal{H}, \boldsymbol{\pi}_i(\boldsymbol{h}) \in \mathcal{X}_i(\boldsymbol{s}^{(\tau)}, \boldsymbol{\pi}_{-i}(\boldsymbol{h}))\},$$

and for any $\mathcal{P}^{\text{sub}} \subseteq \mathcal{P}^{\text{markov}}$, the **feasible subclass policy correspondence**

$$\mathcal{F}_i^{\text{sub}}(\boldsymbol{\pi}_{-i}) \doteq \{\boldsymbol{\pi}_i \in \mathcal{P}_i^{\text{sub}} \mid \forall \boldsymbol{s} \in \mathcal{S}, \boldsymbol{\pi}_i(\boldsymbol{s}) \in \mathcal{X}_i(\boldsymbol{s}, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))\}.$$

Of particular interest is $\mathcal{F}_i^{\text{markov}}(\boldsymbol{\pi}_{-i})$ itself, obtained when $\mathcal{P}^{\text{sub}} = \mathcal{P}^{\text{markov}}$.

Given a policy profile $\boldsymbol{\pi} \in \mathcal{P}$ and a history $\boldsymbol{h} \in \mathcal{H}^\tau$, we define the **discounted history distribution** assuming initial state distribution μ as

$$\nu_\mu^{\boldsymbol{\pi},\tau}(\boldsymbol{h}) = \mu(\boldsymbol{s}^{(0)}) \prod_{t=0}^{\tau-1} \gamma^t \rho(\boldsymbol{s}^{(t+1)} \mid \boldsymbol{s}^{(t)}, \boldsymbol{a}^{(t)}) \mathbb{1}_{\{\boldsymbol{\pi}(\boldsymbol{h}_{:t})\}}(\boldsymbol{a}^{(t)}).$$

Overloading notation, we also define **the set of all realizable trajectories** $\mathcal{H}^{\boldsymbol{\pi}}$ of length τ under policy $\boldsymbol{\pi}$ as $\mathcal{H}^{\boldsymbol{\pi}} \doteq \operatorname{supp}(\nu_\mu^{\boldsymbol{\pi},\tau})$, i.e., the set of all histories that occur with non-zero probability. We then denote by $\nu_\mu^{\boldsymbol{\pi}} \doteq \nu_\mu^{\boldsymbol{\pi},\infty}$, and by $H = \left(S^{(0)}, (A^{(t)}, S^{(t+1)})_{t=0}^{\tau-1}\right)$ any

[2]A **mixed policy** is simply a distribution over pure policies, i.e., an element of $\Delta(\mathcal{P}_i)$. Moreover, any mixed policy can be equivalently represented as a mapping $\boldsymbol{\pi}_i^{\text{mixed}} : \mathcal{H} \to \Delta(\mathcal{A}_i)$ from histories to distributions over actions s.t. at any history $\boldsymbol{h} \in \mathcal{H}$, player i plays action $\boldsymbol{a}_i \sim \boldsymbol{\pi}_i(\boldsymbol{h})$. An analogous definition extends directly to mixed Markov policies as well.

randomly sampled history from $\nu_{\mu}^{\boldsymbol{\pi},\tau}$. Finally, we define the **discounted state-visitation distribution**, again assuming initial state distribution μ, as

$$\delta_{\mu}^{\boldsymbol{\pi}}(\boldsymbol{s}) = \sum_{\tau=0}^{\infty} \int_{\boldsymbol{h} \in \mathcal{H}^{\boldsymbol{\pi}} : \boldsymbol{s}^{(\tau)} = \boldsymbol{s}} \nu_{\mu}^{\boldsymbol{\pi},\tau}(\boldsymbol{h}).$$

For any policy profile $\boldsymbol{\pi} \in \mathcal{P}$, the **state-value function** $\boldsymbol{v}^{\boldsymbol{\pi}} : \mathcal{S} \to \mathbb{R}^n$ and the **action-value function** $\boldsymbol{q}^{\boldsymbol{\pi}} : \mathcal{S} \times \mathcal{A} \to \mathbb{R}^n$ are defined, respectively, as

$$\boldsymbol{v}^{\boldsymbol{\pi}}(\boldsymbol{s}) \doteq \mathop{\mathbb{E}}_{S^{(t+1)} \sim \rho(\cdot | S^{(t)}, A^{(t)})} \left[\sum_{t=0}^{\infty} \boldsymbol{r}(S^{(t)}, A^{(t)}) \mid S^{(0)} = \boldsymbol{s}, A^{(t)} = \boldsymbol{\pi}(S^{(t)}) \right] \tag{12.1}$$

$$\boldsymbol{q}^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{a}) \doteq \mathop{\mathbb{E}}_{S^{(t+1)} \sim \rho(\cdot | S^{(t)}, A^{(t)})} \left[\sum_{t=0}^{\infty} \boldsymbol{r}(S^{(t)}, A^{(t)}) \mid S^{(0)} = \boldsymbol{s}, A^{(0)} = \boldsymbol{a}, A^{(t+1)} = \boldsymbol{\pi}(S^{(t+1)}) \right] . \tag{12.2}$$

Overloading notation, for any arbitrary initial state distribution $\upsilon \in \Delta(\mathcal{S})$ and policy profile $\boldsymbol{\pi}$, we denote by $\boldsymbol{v}^{\boldsymbol{\pi}}(\upsilon) \doteq \mathbb{E}_{S \sim \upsilon}\left[\boldsymbol{v}^{\boldsymbol{\pi}}(S)\right]$.

Finally, the **(expected cumulative) payoff** associated with policy profile $\boldsymbol{\pi} \in \mathcal{P}$ is given by $\boldsymbol{u}(\boldsymbol{\pi}) \doteq \boldsymbol{v}^{\boldsymbol{\pi}}(\mu)$.

12.2 Solution Concepts and Existence

Having defined our model, we now define two natural solution concepts, and establish their existence. Our first solution concept is based on the usual notion of Nash equilibrium (1950b), yet applied to Markov pseudo-games. Our second is based on the notion of subgame-perfect equilibrium in extensive-form games, a strengthening of Nash equilibrium with the additional requirement that an equilibrium be Nash not just at the start of the game, but at all states encountered during play. In the context of stochastic games, such equilibria are called "recursive," or "Markov perfect." Following Bellman (1966) and Arrow and Debreu (1954), we identify natural assumptions that guarantee the existence of equilibrium in (pure) Markov policies, meaning deterministic policies that depend only on the current state, not on the history. When applied to Radner economies, this theorem implies existence of (pure) recursive Radner equilibrium, to our knowledge the first result of its kind.

Definition 12.2.1 [Approximate Generalized Nash Equilibrium]**.**
An ε-**generalized Nash equilibrium (**ε**-GNE)** $\boldsymbol{\pi}^* \in \mathcal{F}(\boldsymbol{\pi}^*)$ is a policy profile s.t. for all states $\boldsymbol{s} \in \mathcal{S}$ and players $i \in [n]$,

$$u_i(\boldsymbol{\pi}^*) \geq \max_{\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}^*_{-i})} u_i(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i}) - \varepsilon \enspace .$$

We call a 0-GNE simply a GNE.

Definition 12.2.2 [Approximate Generalized Markov Perfect Equilibrium]**.**
An ε-**generalized Markov perfect equilibrium (**ε**-GMPE)** $\boldsymbol{\pi}^* \in \mathcal{F}^{\mathrm{markov}}(\boldsymbol{\pi}^*)$ is a Markov policy profile s.t. for all states $\boldsymbol{s} \in \mathcal{S}$ and players $i \in [n]$,

$$v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}) \geq \max_{\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}^*_{-i})} v_i^{(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i})}(\boldsymbol{s}) - \varepsilon \enspace .$$

We call a 0-GMPE simply a GMPE.

As GMPE is a stronger notion than GNE, every ε-GMPE is an ε-GNE.

Our existence result is based on two assumptions, one set of assumptions on the game, and another set of assumptions on the policy subclass for which existence is sought. The following assumption on the Markov pseudo-game ensures that the Markov pseudo-game is "continuous" and "concave".

Assumption 12.2.1 [Concave Markov Pseudo-game]**.**

For all players $i \in [n]$, assume

1. $\mathcal{A}_i$ is convex
2. $\mathcal{X}_i(\boldsymbol{s}, \cdot)$ is upper- and lower-hemicontinuous, for all $\boldsymbol{s} \in \mathcal{S}$
3. $\mathcal{X}_i(\boldsymbol{s}, \boldsymbol{a}_{-i})$ is non-empty, convex, and compact, for all $\boldsymbol{s} \in \mathcal{S}$ and $\boldsymbol{a}_{-i} \in \mathcal{A}_{-i}$
4. For any policy $\boldsymbol{\pi} \in \mathcal{P}$, $\boldsymbol{a}_i \mapsto q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{a}_i, \boldsymbol{a}_{-i})$ is continuous and concave over $\mathcal{X}_i(\boldsymbol{s}, \boldsymbol{a}_{-i})$, for all $\boldsymbol{s} \in \mathcal{S}$ and $\boldsymbol{a}_{-i} \in \mathcal{A}_{-i}$

The following assumption of the policy subclass ensures that the policy subclass is expressive and well-behaved enough for it to represent a GMPE. We note that the set of Markov policies by definition satisfies the following assumption.

Assumption 12.2.2 [Policy Class]**.**

Given $\mathcal{P}^{\text{sub}} \subseteq \mathcal{P}^{\text{markov}}$, assume:

1. $\mathcal{P}^{\text{sub}}$ is non-empty, compact, and convex
2. (Closure under policy improvement): For each $\boldsymbol{\pi} \in \mathcal{P}^{\text{sub}}$, there exists $\boldsymbol{\pi}^+ \in \mathcal{P}^{\text{sub}}$ s.t. $q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{\pi}_i^+(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s})) = \max_{\boldsymbol{\pi}_i' \in \mathcal{F}(\boldsymbol{\pi}_{-i})} q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{\pi}_i'(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s}))$, for all $i \in [n]$ and $\boldsymbol{s} \in \mathcal{S}$

Assumption 2, introduced as Condition 1 in Bhandari and Russo (2019), ensures that the policy class under consideration (e.g., $\mathcal{P}^{\text{sub}} \subseteq \mathcal{P}^{\text{markov}}$) is expressive enough to include best responses. With the above assumptions in hand, we can prove the existence of a GMPE using the Kakutani-Glicksberg fixed point theorem (Glicksberg, 1952) (Theorem 2.4.1, Chapter 2)

Theorem 12.2.1.

Let $\mathcal{M}$ be a Markov pseudo-game for which Assumption 12.2.1 holds, and let $\mathcal{P}^{\text{sub}} \subseteq \mathcal{P}^{\text{markov}}$ be a subspace of Markov policy profiles that satisfies Assumption 12.2.2. Then, there exists a policy $\boldsymbol{\pi}^* \in \mathcal{P}^{\text{sub}}$ such that $\boldsymbol{\pi}^*$ is an GMPE of $\mathcal{M}$.

With solution concepts and their existence, we next turn our attention to computation.

12.3 Merit Function Minimization for Generalized Markov Perfect Equilibrium

Our approach to computing a GMPE in a Markov pseudo-game $\mathcal{M}$ is to minimize a **merit function** associated with $\mathcal{M}$, i.e., a function whose minima coincides with the pseudo-game's GMPE. Our choice of merit function, a common one in game theory, is **exploitability** $\varphi : \mathcal{P} \to \mathbb{R}_+$, defined as $\varphi(\boldsymbol{\pi}) \doteq \sum_{i\in[n]} \left[\max_{\boldsymbol{\pi}'_i \in \mathcal{F}_i^{\text{markov}}(\boldsymbol{\pi}_{-i})} u_i(\boldsymbol{\pi}'_i, \boldsymbol{\pi}_{-i}) - u_i(\boldsymbol{\pi})\right]$. In words, exploitability is the sum of the players' maximal unilateral payoff deviations.

Exploitability, however, is a merit function for GNE, *not* GMPE; **state exploitability**, $\phi(\boldsymbol{s}, \boldsymbol{\pi}) = \sum_{i\in[n]}[\max_{\boldsymbol{\pi}'_i \in \mathcal{F}_i^{\text{markov}}(\boldsymbol{\pi}_{-i})} v_i^{(\boldsymbol{\pi}'_i, \boldsymbol{\pi}_{-i})}(\boldsymbol{s}) - v_i^{\boldsymbol{\pi}}(\boldsymbol{s})]$ at all states $\boldsymbol{s} \in \mathcal{S}$, is a merit function for GMPE. Nevertheless, as we show in the sequel, for a large class of Markov pseudo-games, namely those with a bounded best-response mismatch coefficient (see Section 12.3.3), the set of Markov policies that minimize exploitability equals the set of GMPE, making our approach a sensible one.

We are not out of the woods yet, however, as exploitability is non-convex in general, even in one-shot finite games (Nash, 1950a). Although Markov pseudo-games can afford a convex exploitability (see, for instance (Flam and Ruszczynski, 1994)), it is unlikely that all do, as GNE computation is PPAD-hard (Chen et al., 2009; Daskalakis et al., 2009). Accordingly, we instead set our sights on computing a **stationary point** of the exploitability, i.e., a policy profile $\boldsymbol{\pi}^* \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)$ s.t. for any other policy $\boldsymbol{\pi} \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)$, it holds that $\min_{\boldsymbol{h}\in\mathcal{D}\varphi(\boldsymbol{\pi}^*)} \langle \boldsymbol{h}, \boldsymbol{\pi}^* - \boldsymbol{\pi}\rangle \leq 0$.[3] Such a point satisfies the necessary conditions of a GMPE.

In this paper, we study Markov pseudo-games with possibly continuous state and action spaces. As such, we can only hope to compute an *approximate* stationary point of the exploitability in finite time. Defining a notion of approximate stationarity for exploitability

[3]While we provide a definition of a(n approximate) stationary point for expositional purposes at present, an observant reader might have noticed the exploitability φ is a mapping from a function space to the positive reals, and its Frêchet (sub)derivative is ill-specified without a clear definition of the normed vector space of policies on which exploitability is defined. Further, even when clearly specified, such a (sub)derivative might not exist. The precise meaning of a derivative of the exploitability and its stationary points will be introduced more rigorously once we have suitably parameterized the policy spaces.

is, however, a challenge, because exploitability is non-differentiable in general (once again, even in one-shot finite games).

Given an approximation parameter $\varepsilon \geq 0$, a natural definition of an ε-stationary point might be a policy profile $\boldsymbol{\pi}^* \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)$ s.t. for any other policy $\boldsymbol{\pi} \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)$, it holds that $\min_{\boldsymbol{h} \in \mathcal{D}\varphi(\boldsymbol{\pi}^*)} \langle \boldsymbol{h}, \boldsymbol{\pi}^* - \boldsymbol{\pi} \rangle \leq \varepsilon$. Exploitability is not necessarily Lipschitz-smooth, however, so in general it may not be possible to compute an ε-stationary point in $\text{poly}(1/\varepsilon)$ evaluations of the (sub)gradient of the exploitability.[4]

To address, this challenge, a common approach in the optimization literature (see, for instance Appendix H, Definition 19 of Liu et al. (2021)) is to consider an alternative definition known as (ε, δ)-stationarity. Given approximation parameters $\varepsilon, \delta \geq 0$, an (ε, δ)-stationary point of the exploitability is a policy profile $\boldsymbol{\pi}^* \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)$ for which there exists a δ-close policy $\boldsymbol{\pi}^\dagger \in \mathcal{P}$ with $\|\boldsymbol{\pi}^\dagger - \boldsymbol{\pi}^*\| \leq \delta$ s.t. for any other policy $\boldsymbol{\pi} \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^\dagger)$, it holds that $\min_{\boldsymbol{h} \in \mathcal{D}\varphi(\boldsymbol{\pi}^\dagger)} \langle \boldsymbol{h}, \boldsymbol{\pi}^\dagger - \boldsymbol{\pi} \rangle \leq \varepsilon$. The exploitability minimization method we introduce can compute such an approximate stationary point in polynomial time. Furthermore, asymptotically, our method is guaranteed to converge to an exact stationary point of the exploitability.

More precisely, following Goktas and Greenwald (2022b), who minimize exploitability to solve for variational equilibria in (one-shot) pseudo-games, we first formulate our problem as the quasi-optimization problem of minimizing exploitability,[5] and then transform this problem into a coupled min-max optimization (i.e., a two-player zero-sum game) whose objective is cumulative regret, rather than the potentially ill-behaved exploitability. Under suitable parametrization, such problems are amenable to polynomial-time solutions via si-

[4]To see this, consider the convex minimization problem $\min_{x \in \mathbb{R}} f(x) = |x|$. The minimum of this optimization occurs at $x = 0$, which is a stationary point since a (sub)derivative of f at $x = 0$ is 0. However, for $x < 0$, we have $\frac{\partial f(x)}{\partial x} = -1$, and for $x > 0$, we have $\frac{\partial f(x)}{\partial x} = 1$. Hence, any $x \in \mathbb{R} \setminus \{0\}$ can at best be a 1-stationary point, i.e., $\left|\frac{\partial f(x)}{\partial x}\right| = 1$. Hence, for this optimization problem, it is not even possible to guarantee the existence of an ε-stationary point distinct from $x = 0$, assuming $\varepsilon \in (0, 1)$, let alone the computation of an ε-stationary point x^* s.t. $\left|\frac{\partial f(x^*)}{\partial x}\right| \leq \varepsilon$.

[5]Here, "quasi" refers to the fact that a solution to this problem is both a minimizer of exploitability and a fixed point of an operator, such as $\mathcal{F}$ or $\mathcal{F}^{\text{markov}}$.

multaneous gradient descent ascent (Arrow et al., 1958), assuming the objective is Lipschitz smooth in both players' decision variables and gradient dominated in the inner player's. We thus formulate the requisite assumptions to ensure these properties hold of cumulative regret in our game, which in turn allows us to show that **two time scale simultaneous stochastic gradient descent ascent (TTSSGDA)** converges to an $(\varepsilon, O(\varepsilon))$-stationary point of the exploitability in $\text{poly}(1/\varepsilon)$ gradient steps.

12.3.1 Exploitability Minimization

Given a Markov pseudo-game $\mathcal{M}$ and two policy profiles $\boldsymbol{\pi}, \boldsymbol{\pi}' \in \mathcal{P}$, we define the **state cumulative regret** at state $\boldsymbol{s} \in \mathcal{S}$ as $\psi(\boldsymbol{s}, \boldsymbol{\pi}, \boldsymbol{\pi}') = \sum_{i \in [n]} \left[v_i^{(\boldsymbol{\pi}'_i, \boldsymbol{\pi}_{-i})}(\boldsymbol{s}) - v_i^{\boldsymbol{\pi}}(\boldsymbol{s}) \right]$; the **expected cumulative regret** for any initial state distribution $\upsilon \in \Delta(\mathcal{S})$ as $\psi(\upsilon, \boldsymbol{\pi}, \boldsymbol{\pi}') = \mathbb{E}_{\boldsymbol{s} \sim \upsilon} [\psi(\boldsymbol{s}, \boldsymbol{\pi}, \boldsymbol{\pi}')]$, and the **cumulative regret** as $\Psi(\boldsymbol{\pi}, \boldsymbol{\pi}') = \psi(\mu, \boldsymbol{\pi}, \boldsymbol{\pi}')$. Additionally, we define the **state exploitability** of a policy profile $\boldsymbol{\pi}$ at state $\boldsymbol{s} \in \mathcal{S}$ as $\phi(\boldsymbol{s}, \boldsymbol{\pi}) = \sum_{i \in [n]} \max_{\boldsymbol{\pi}'_i \in \mathcal{F}_i^{\text{markov}}(\boldsymbol{\pi}_{-i})} v_i^{(\boldsymbol{\pi}'_i, \boldsymbol{\pi}_{-i})}(\boldsymbol{s}) - v_i^{\boldsymbol{\pi}}(\boldsymbol{s});$, the **expected exploitability** of a policy profile $\boldsymbol{\pi}$ for any initial state distribution $\upsilon \in \Delta(\mathcal{S})$ as $\phi(\upsilon, \boldsymbol{\pi}) = \mathbb{E}_{\boldsymbol{s} \sim \upsilon} [\phi(\boldsymbol{s}, \boldsymbol{\pi})]$, and **exploitability** as $\varphi(\boldsymbol{\pi}) = \sum_{i \in [n]} \max_{\boldsymbol{\pi}'_i \in \mathcal{F}_i^{\text{markov}}(\boldsymbol{\pi}_{-i})} u_i(\boldsymbol{\pi}'_i, \boldsymbol{\pi}_{-i})$.

In the above, we restrict our attention to the subclass $\mathcal{P}^{\text{markov}} \subseteq \mathcal{P}$ of (pure) Markov policies. This restriction is without loss of generality, because finding an optimal policy that maximizes a state-value or payoff function, while the other players' policies remain fixed, reduces to solving a Markov decision process (MDP), and an optimal (possibly history-dependent) policy in an MDP is guaranteed to exist in the space of (pure) Markov policies under very mild continuity and compactness assumptions (Puterman, 2014). Indeed, the next lemma justifies this restriction.

Lemma 12.3.1.

Given a Markov pseudo-game $\mathcal{M}$ for which Assumption 12.2.1 holds, a Markov policy profile $\boldsymbol{\pi}^* \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)$ is a GMPE if and only if $\phi(\boldsymbol{s}, \boldsymbol{\pi}^*) = 0$, for all states $\boldsymbol{s} \in \mathcal{S}$. Similarly, a policy profile $\boldsymbol{\pi}^* \in \mathcal{F}(\boldsymbol{\pi}^*)$ is an GNE if and only if $\varphi(\boldsymbol{\pi}^*) = 0$.

This lemma tells us that we can reformulate the problem of computing a GMPE as the quasi-minimization problem of minimizing state exploitability, i.e., $\min_{\boldsymbol{\pi} \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi})} \phi(\boldsymbol{s}, \boldsymbol{\pi})$, at all states $\boldsymbol{s} \in \mathcal{S}$ simultaneously. The same is true of computing a GNE and exploitability.

This straightforward reformulation of GMPE (resp. GNE) in terms of state exploitability (resp. exploitability) does not immediately lend itself to computation, as exploitability minimization is non-trivial, because exploitability is neither convex nor differentiable in general. Following Goktas and Greenwald (2022b), we can reformulate these problems yet again, this time as coupled quasi-min-max optimization problems (Wald, 1945). We proceed to do so now; however, we restrict our attention to exploitability, and hence GNE, knowing that we will later show that minimizing exploitability suffices to minimize state exploitability, and thereby find GMPE.

Observation 12.3.1.

Given a Markov pseudo-game $\mathcal{M}$,

$$\min_{\boldsymbol{\pi} \in \mathcal{F}(\boldsymbol{\pi})} \varphi(\boldsymbol{\pi}) = \min_{\boldsymbol{\pi} \in \mathcal{F}(\boldsymbol{\pi})} \max_{\boldsymbol{\pi}' \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi})} \Psi(\boldsymbol{\pi}, \boldsymbol{\pi}') \ . \tag{12.3}$$

While the above observation makes progress towards our goal of reformulating exploitability minimization in a tractable manner, the problem remains challenging to solve for two reasons: first, a fixed point computation is required to solve the outer player's minimization problem; second, the inner player's policy space depends on the choice of outer policy. We overcome these difficulties by choosing suitable policy parameterizations.

12.3.2 Policy Parameterization

In a coupled min-max optimization problem, any solution to the inner player's maximization problem is implicitly parameterized by the outer player's decision. We restructure the jointly feasible Markov policy class to represent this dependence explicitly.

Define the class of **dependent policies** $\mathcal{R} \doteq \{\boldsymbol{\rho} : \mathcal{S} \times \mathcal{A} \to \mathcal{A} \mid \forall (\boldsymbol{s}, \boldsymbol{a}) \in \mathcal{S} \times \mathcal{A},\ \boldsymbol{\rho}(\boldsymbol{s}, \boldsymbol{a}) \in \mathcal{X}(\boldsymbol{s}, \boldsymbol{a})\} = \bigtimes_{i \in [n]} \{\boldsymbol{\rho}_i : \mathcal{S} \times \mathcal{A}_i \to \mathcal{A}_{-i} \mid \forall (\boldsymbol{s}, \boldsymbol{a}_{-i}) \in \mathcal{S} \times \mathcal{A}_{-i},\ \boldsymbol{\rho}_i(\boldsymbol{s}, \boldsymbol{a}_{-i}) \in \mathcal{X}_i(\boldsymbol{s}, \boldsymbol{a}_{-i})\}$. With this definition in hand we arrive at an *uncoupled* quasi-min-max optimization problem:

Lemma 12.3.2.

Given a Markov pseudo-game $\mathcal{M}$,

$$\min_{\boldsymbol{\pi} \in \mathcal{F}(\boldsymbol{\pi})} \max_{\boldsymbol{\pi}' \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi})} \Psi(\boldsymbol{\pi}, \boldsymbol{\pi}') = \min_{\boldsymbol{\pi} \in \mathcal{F}(\boldsymbol{\pi})} \max_{\boldsymbol{\rho} \in \mathcal{R}} \Psi(\boldsymbol{\pi}, \boldsymbol{\rho}(\cdot, \boldsymbol{\pi}(\cdot))) \ . \tag{12.4}$$

It can be expensive to represent the aforementioned dependence in policies explicitly. This situation can be naturally rectified, however, by a suitable policy parameterization. A suitable policy parameterization can also allow us to represent the set of fixed points s.t. $\boldsymbol{\pi} \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi})$ more efficiently in practice (Goktas et al., 2023a).

Define a **parameterization scheme** $(\boldsymbol{\pi}, \boldsymbol{\rho}, \mathbb{R}^{\Omega}, \mathbb{R}^{\Sigma})$ as comprising a unconstrained parameter space $\mathbb{R}^{\Omega}$ and parametric policy profile function $\boldsymbol{\pi} : \mathcal{S} \times \mathbb{R}^{\Omega} \to \mathcal{A}$ for the outer player, and an unconstrained parameter space $\mathbb{R}^{\Sigma}$ and parametric policy profile function $\boldsymbol{\rho} : \mathcal{S} \times \mathcal{A} \times \mathbb{R}^{\Sigma} \to \mathcal{A}$ for the inner player. Given such a scheme, we restrict the players' policies to be parameterized: i.e., the outer player's space of policies $\mathcal{P}^{\mathbb{R}^{\Omega}} = \{\boldsymbol{\pi} : \mathcal{S} \times \mathbb{R}^{\Omega} \to \mathcal{A} \mid \boldsymbol{\omega} \in \mathbb{R}^{\Omega}\} \subseteq \mathcal{P}^{\text{markov}}$, while the inner player's space of policies $\mathcal{R}^{\mathbb{R}^{\Sigma}} = \{\boldsymbol{\rho} : \mathcal{S} \times \mathcal{A} \times \mathbb{R}^{\Sigma} \to \mathcal{A} \mid \boldsymbol{\sigma} \in \mathbb{R}^{\Sigma}\}$. Using these parametrization, we redefine $\boldsymbol{v}^{\boldsymbol{\omega}} \doteq \boldsymbol{v}^{\boldsymbol{\pi}(\cdot;\boldsymbol{\omega})}$, $\boldsymbol{q}^{\boldsymbol{\omega}} \doteq \boldsymbol{q}^{\boldsymbol{\pi}(\cdot;\boldsymbol{\omega})}$, $\boldsymbol{u}(\boldsymbol{\omega}) = \boldsymbol{u}(\boldsymbol{\pi}(\cdot;\boldsymbol{\omega}))$, and $\nu_{\mu}^{\boldsymbol{\omega}} = \nu_{\mu}^{\boldsymbol{\pi}(\cdot;\boldsymbol{\omega})}$; $\boldsymbol{v}^{\boldsymbol{\sigma}(\boldsymbol{\omega})} \doteq \boldsymbol{v}^{\boldsymbol{\rho}(\cdot,\boldsymbol{\pi}(\cdot;\boldsymbol{\omega});\boldsymbol{\sigma})}$; $\boldsymbol{q}^{\boldsymbol{\sigma}(\boldsymbol{\omega})} \doteq \boldsymbol{q}^{\boldsymbol{\rho}(\cdot,\boldsymbol{\pi}(\cdot;\boldsymbol{\omega});\boldsymbol{\sigma})}$; $\boldsymbol{u}(\boldsymbol{\sigma}(\boldsymbol{\omega})) = \boldsymbol{u}(\boldsymbol{\rho}(\cdot, \boldsymbol{\pi}(\cdot;\boldsymbol{\omega});\boldsymbol{\sigma}))$; $\nu_{\mu}^{\boldsymbol{\sigma}(\boldsymbol{\omega})} = \nu_{\mu}^{\boldsymbol{\rho}(\cdot,\boldsymbol{\pi}(\cdot;\boldsymbol{\omega});\boldsymbol{\sigma})}$. With these definitions in place, we make the following assumption on our parametrization.

Assumption 12.3.1 [Parametrization for Min-Max Optimization]**.**
Given a Markov pseudo-game $\mathcal{M}$ and a parameterization scheme $(\boldsymbol{\pi}, \boldsymbol{\rho}, \mathbb{R}^{\Omega}, \mathbb{R}^{\Sigma})$, assume:

1. for all $\boldsymbol{\omega} \in \mathbb{R}^{\Omega}$, $\boldsymbol{\pi}(\boldsymbol{s};\boldsymbol{\omega}) \in \mathcal{X}(\boldsymbol{s}, \boldsymbol{\pi}(\boldsymbol{s};\boldsymbol{\omega}))$, for all $\boldsymbol{s} \in \mathcal{S}$
2. for all $\boldsymbol{\sigma} \in \mathbb{R}^{\Sigma}$, $\boldsymbol{\rho}(\boldsymbol{s}, \boldsymbol{a};\boldsymbol{\sigma}) \in \mathcal{X}(\boldsymbol{s}, \boldsymbol{a})$, for all $(\boldsymbol{s}, \boldsymbol{a}) \in \mathcal{S} \times \mathcal{A}$

Assuming a policy parameterization scheme that satisfies Assumption 12.3.1, we restate our goal, state exploitability minimization, one last time as the following min-max optimization problem:

$$\min_{\boldsymbol{\omega} \in \mathbb{R}^{\Omega}} \max_{\boldsymbol{\sigma} \in \mathbb{R}^{\Sigma}} \Psi(\boldsymbol{\omega}, \boldsymbol{\sigma}) \doteq \Psi(\boldsymbol{\pi}(\cdot;\boldsymbol{\omega}), \boldsymbol{\rho}(\cdot, \boldsymbol{\pi}(\cdot;\boldsymbol{\omega});\boldsymbol{\sigma})) \ . \tag{12.5}$$

Now, given unconstrained parameter space, we are able to simplify our definition of approximate stationary point and obtain our target definition.

Definition 12.3.1.

Given $\varepsilon, \delta \geq 0$, a (ε, δ)**-stationary point of the exploitability** is a policy parameter $\boldsymbol{\omega}^* \in \mathbb{R}^\Omega$ for which there exists a δ-close policy parameter $\boldsymbol{\omega}^\dagger \in \mathbb{R}^\Omega$ with $\|\boldsymbol{\omega}^* - \boldsymbol{\omega}^\dagger\| \leq \delta$ s.t. $\min_{\boldsymbol{h} \in \mathcal{D}\varphi(\boldsymbol{\omega}^\dagger)} \|\boldsymbol{h}\| \leq \varepsilon$.

12.3.3 State Exploitability Minimization

Returning to our stated objective, namely *state* exploitability minimization, we turn our attention to obtaining a tractable characterization of this goal. Specifically, we argue that it suffices to minimize exploitability, rather than state exploitability, as any policy profile that is a stationary point of exploitability is also a stationary point of state exploitability across all states simultaneously, under suitable assumptions.

Our first lemma states that a stationary point of the exploitability is almost surely also a stationary point of the state exploitability at all states. Moreover, if the initial state distribution has full support, then any (ε, δ)-stationary point of the exploitability can be converted into an $(\varepsilon/\alpha, \delta)$-stationary point of the *state* exploitability, with probability at least $1 - \alpha$.

Lemma 12.3.3.

Given a Markov pseudo-game $\mathcal{M}$, for $\boldsymbol{\omega} \in \mathbb{R}^\Omega$, suppose that $\phi(\boldsymbol{s}, \cdot)$ is differentiable at $\boldsymbol{\omega}$ for all $\boldsymbol{s} \in \mathcal{S}$. If $\|\nabla_{\boldsymbol{\omega}} \varphi(\boldsymbol{\omega})\| = 0$, then, for all states $\boldsymbol{s} \in \mathcal{S}$, $\|\nabla_{\boldsymbol{\omega}} \phi(\boldsymbol{s}, \boldsymbol{\omega})\| = 0$ μ-almost surely, i.e., $\mu(\{\boldsymbol{s} \in \mathcal{S} \mid \|\nabla_{\boldsymbol{\omega}} \phi(\boldsymbol{s}, \boldsymbol{\omega})\| = 0\}) = 1$. Moreover, for any $\varepsilon > 0$ and $\delta \in [0, 1]$, if $\operatorname{supp}(\mu) = \mathcal{S}$ and $\|\nabla_{\boldsymbol{\omega}} \varphi(\boldsymbol{\omega})\| \leq \varepsilon$, then with probability at least $1 - \delta$, $\|\nabla_{\boldsymbol{\omega}} \phi(\boldsymbol{s}, \boldsymbol{\omega})\| \leq \varepsilon/\delta$.

In fact, we can strengthen this probabilistic equivalence to a deterministic one by restricting our attention to Markov pseudo-games with bounded best-response mismatch coefficients. Our best-response mismatch coefficient generalizes the minimax mismatch coefficient in two-player settings (Daskalakis et al., 2020a) and the distribution mismatch coefficient in single-agent settings (Agarwal et al., 2020).

Definition 12.3.2 [Best-Response Mismatch Coefficient]**.**

Given $\mathcal{M}$ with initial state distribution μ and alternative state distribution $\upsilon \in \Delta(\mathcal{S})$, and letting $\Phi_i(\boldsymbol{\pi}_{-i}) \doteq \arg\max_{\boldsymbol{\pi}_i' \in \mathcal{F}_i^{\text{markov}}(\boldsymbol{\pi}_{-i})} u_i(\boldsymbol{\pi}_i', \boldsymbol{\pi}_{-i})$ denote the set of best response policies for player i when the other players play policy profile $\boldsymbol{\pi}_{-i}$, we define the **best-response mismatch coefficient** for policy profile $\boldsymbol{\pi}$ as

$$C_{br}(\boldsymbol{\pi}, \mu, \upsilon) \doteq \max_{i \in [n]} \max_{\boldsymbol{\pi}_i' \in \Phi_i(\boldsymbol{\pi}_{-i})} \left(\frac{1}{1-\gamma}\right)^2 \left\| \frac{\delta_\upsilon^{(\boldsymbol{\pi}_i', \boldsymbol{\pi}_{-i})}}{\mu} \right\|_\infty \left\| \frac{\delta_\upsilon^{\boldsymbol{\pi}}}{\mu} \right\|_\infty .$$

Lemma 12.3.4.

Let $\mathcal{M}$ be a Markov pseudo-game with initial state distribution μ. Given policy parameter $\boldsymbol{\omega} \in \mathbb{R}^\Omega$ and arbitrary state distribution $\upsilon \in \Delta(\mathcal{S})$, suppose that both $\phi(\mu, \cdot)$ and $\phi(\upsilon, \cdot)$ are differentiable at $\boldsymbol{\omega}$, then we have: $\|\nabla \phi(\upsilon, \boldsymbol{\omega})\| \leq C_{br}(\boldsymbol{\pi}(\cdot; \boldsymbol{\omega}^*), \mu, \upsilon) \|\nabla \varphi(\boldsymbol{\omega})\|$.

Once again, Lemma 12.3.3 states that any approximate stationary point of exploitability is also an approximate stationary point of state exploitability with high probability, while Lemma 12.3.4 upper bounds state exploitability in terms of exploitability, when the best-response mismatch coefficient is bounded. Together, these two lemmas imply that finding a policy profile this is a stationary point of exploitability is sufficient for find a policy profile this is a stationary point of state exploitability, and hence one that satisfies the necessary conditions of a GMPE.

12.3.4 Algorithmic Assumptions

We are nearly ready to describe our reinforcement learning algorithm for computing a stationary point of Equation (12.5), and thereby finding a policy profile that satisfies the necessary conditions of a GMPE. As Equation (12.5) is a two-player zero-sum game, our method is a variant of simultaneous gradient descent ascent (GDA) (Arrow et al., 1958), meaning it adjusts its parameters based on first-order information until it reaches a (first-order) stationary point. Polynomial-time convergence of GDA typically requires that the objective be Lipschitz smooth in both decision variables, and gradient dominated in the inner one, which in our application, translates to the cumulative regret $\Psi(\boldsymbol{\omega}, \boldsymbol{\sigma})$ being Lipschitz smooth in $(\boldsymbol{\omega}, \boldsymbol{\sigma})$ and gradient dominated in $\boldsymbol{\sigma}$. These conditions are ensured, under the following assumptions on the Markov pseudo-game.

Assumption 12.3.2 [Lipschitz Smooth Payoffs]**.**

Given a Markov pseudo-game $\mathcal{M}$ and a parameterization scheme $(\boldsymbol{\pi}, \boldsymbol{\rho}, \mathbb{R}^\Omega, \mathbb{R}^\Sigma)$, assume:

1. $\mathbb{R}^\Omega$ and $\mathbb{R}^\Sigma$ are non-empty, compact, and convex
2. $\boldsymbol{\omega} \mapsto \boldsymbol{\pi}(\boldsymbol{s}; \boldsymbol{\omega})$ is twice continuously differentiable, for all $\boldsymbol{s} \in \mathcal{S}$, and $\boldsymbol{\sigma} \mapsto \boldsymbol{\rho}(\boldsymbol{s}, \boldsymbol{a}; \boldsymbol{\sigma})$ is twice continuously differentiable, for all $(\boldsymbol{s}, \boldsymbol{a}) \in \mathcal{S} \times \mathcal{A}$
3. $\boldsymbol{a} \mapsto \boldsymbol{r}(\boldsymbol{s}, \boldsymbol{a})$ is twice continuously differentiable, for all $\boldsymbol{s} \in \mathcal{S}$
4. $\boldsymbol{a} \mapsto \rho(\boldsymbol{s}' \mid \boldsymbol{s}, \boldsymbol{a})$ is twice continuously differentiable, for all $\boldsymbol{s}, \boldsymbol{s}' \in \mathcal{S}$.

Assumption 12.3.3 [Gradient Dominance Conditions]**.**

Given a Markov pseudo-game $\mathcal{M}$ together with a parameterization scheme $(\boldsymbol{\pi}, \boldsymbol{\rho}, \mathbb{R}^\Omega, \mathbb{R}^\Sigma)$, assume:

1. (Closure under policy improvement): For each $\boldsymbol{\omega} \in \mathbb{R}^\Omega$, there exists $\boldsymbol{\sigma} \in \mathbb{R}^\Sigma$ s.t. $q_i^{\boldsymbol{\omega}}(\boldsymbol{s}, \boldsymbol{\rho}_i(\boldsymbol{s}, \boldsymbol{\pi}(\boldsymbol{s}; \boldsymbol{\omega}); \boldsymbol{\sigma}), \boldsymbol{\pi}_{-i}(\boldsymbol{s}; \boldsymbol{\omega})) = \max_{\boldsymbol{\pi}_i' \in \mathcal{F}_i(\boldsymbol{\pi}(\cdot; \boldsymbol{\omega}))} q_i^{\boldsymbol{\omega}}(\boldsymbol{s}, \boldsymbol{\pi}_i'(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s}; \boldsymbol{\omega}))$ for all $i \in [n]$, $\boldsymbol{s} \in \mathcal{S}$.

2. (Concavity of cumulative regret) $\boldsymbol{\sigma} \mapsto q_i^{\boldsymbol{\omega}'}(\boldsymbol{s}, \boldsymbol{\rho}_i(\boldsymbol{s}, \boldsymbol{\pi}_{-i}(\boldsymbol{s}; \boldsymbol{\omega}); \boldsymbol{\sigma}), \boldsymbol{\pi}_{-i}(\boldsymbol{s}; \boldsymbol{\omega}))$ is concave, for all $\boldsymbol{s} \in \mathcal{S}$ and $\boldsymbol{\omega}, \boldsymbol{\omega}' \in \mathbb{R}^\Omega$.

12.3.5 Algorithm and Convergence

Finally, we present our algorithm for finding an approximate stationary point of exploitability, and thus state exploitability. The algorithm we use is two time-scale stochastic simultaneous gradient descent-ascent (TTSSGDA), first analyzed by Lin et al. (2020); Daskalakis et al. (2020a), for which we prove best-iterate convergence to an $(\varepsilon, O(\varepsilon))$-stationary point of exploitability after taking $\mathrm{poly}(1/\varepsilon)$ gradient steps under Assumptions 12.3.2 and 12.3.3.

Algorithm 11 Two time-scale simultaneous SGDA (TTSSGDA)

Inputs: $\mathcal{M}, (\boldsymbol{\pi}, \boldsymbol{\rho}, \mathbb{R}^{\Omega}, \mathbb{R}^{\Sigma}), \eta_{\boldsymbol{\omega}}, \eta_{\boldsymbol{\sigma}}, \boldsymbol{\omega}^{(0)}, \boldsymbol{\sigma}^{(0)}, T$

Outputs: $(\boldsymbol{\omega}^{(t)}, \boldsymbol{\sigma}^{(t)})_{t=0}^{T}$

1: Build gradient estimator $\widehat{G}$ associated with $\mathcal{M}$

2: **for** $t = 0, \ldots, T-1$ **do**

3: $\quad \boldsymbol{h} \sim \nu^{\boldsymbol{\omega}}, \boldsymbol{h}' \sim \bigtimes_{i \in [n]} \nu^{(\boldsymbol{\sigma}_i(\boldsymbol{\omega}_{-i}), \boldsymbol{\omega}_{-i})}$

4: $\quad \boldsymbol{\omega}^{(t+1)} \leftarrow \boldsymbol{\omega}^{(t)} - \eta_{\boldsymbol{\omega}} \widehat{G_{\boldsymbol{\omega}}}(\boldsymbol{\omega}^{(t)}, \boldsymbol{\sigma}^{(t)}; \boldsymbol{h}, \boldsymbol{h}')$

5: $\quad \boldsymbol{\sigma}^{(t+1)} \leftarrow \boldsymbol{\sigma}^{(t)} + \eta_{\boldsymbol{\sigma}} \widehat{G_{\boldsymbol{\sigma}}}(\boldsymbol{\omega}^{(t)}, \boldsymbol{\sigma}^{(t)}; \boldsymbol{h}, \boldsymbol{h}')$

6: **return** $(\boldsymbol{\omega}^{(t)}, \boldsymbol{\sigma}^{(t)})_{t=0}^{T}$

Recall that Assumption 12.3.2 guarantees Lipschitz smoothness w.r.t. to both $\boldsymbol{\omega}$ and $\boldsymbol{\sigma}$, while Assumption 12.3.3 guarantees gradient dominance w.r.t $\boldsymbol{\sigma}$. As the gradient of cumulative regret involves an expectation over histories, we assume that we can simulate trajectories of play $\boldsymbol{h} \sim \nu_{\mu}^{\boldsymbol{\pi}}$ according to the history distribution $\nu_{\mu}^{\boldsymbol{\pi}}$, for any policy profile $\boldsymbol{\pi}$, and that doing so provides both value and gradient information for the rewards and transition probabilities along simulated trajectories. That is, we rely on a differentiable game simulator (see, for instance Suh et al. (2022)), meaning a stochastic first-order oracle that returns the gradients of the rewards $\boldsymbol{r}$ and transition probabilities ρ, which we query to estimate deviation payoffs, and ultimately cumulative regrets.

Under this assumption, we estimate these values using realized trajectories from the history distribution $\boldsymbol{h} \sim \nu_{\mu}^{\boldsymbol{\omega}}$ induced by the outer player's policy, and the deviation history distribution $\boldsymbol{h}^{\boldsymbol{\sigma}} \sim \bigtimes_{i \in [n]} \nu_{\mu}^{(\boldsymbol{\sigma}_i(\boldsymbol{\omega}_{-i}), \boldsymbol{\omega}_{-i})}$ induced by the inner player's policy. More specifically, for all policies $\boldsymbol{\pi} \in \mathcal{P}^{\mathrm{markov}}$ and histories $\boldsymbol{h} \in \mathcal{H}^{\tau}$, the **payoff estimator** for player $i \in [n]$ is given by:

$$\widehat{u_i}(\boldsymbol{\pi}; \boldsymbol{h}) \doteq \sum_{t=0}^{\tau-1} \mu(\boldsymbol{s}^{(0)}) r_i(\boldsymbol{s}^{(t)}, \boldsymbol{\pi}'(\boldsymbol{s}^{(t)})) \prod_{k=0}^{t-1} \gamma^k \rho(\boldsymbol{s}^{(k+1)} \mid \boldsymbol{s}^{(k)}, (\boldsymbol{s}^{(k)})) \ .$$

Furthermore, for all $\boldsymbol{\omega} \in \mathbb{R}^{\Omega}$, $\boldsymbol{\sigma} \in \mathbb{R}^{\Sigma}$, $\boldsymbol{h} \sim \nu_{\mu}^{\boldsymbol{\omega}}$, and $\boldsymbol{h}^{\boldsymbol{\sigma}} = (\boldsymbol{h}_1^{\boldsymbol{\sigma}}, \cdots, \boldsymbol{h}_n^{\boldsymbol{\sigma}}) \sim \bigtimes_{i \in [n]} \nu_{\mu}^{(\boldsymbol{\sigma}_i(\boldsymbol{\omega}_{-i}), \boldsymbol{\omega}_i)}$, the **cumulative regret estimator** is given by $\widehat{\Psi}(\boldsymbol{\omega}, \boldsymbol{\sigma}; \boldsymbol{h}, \boldsymbol{h}') \doteq \sum_{i \in [n]} \widehat{u}_i(\boldsymbol{\rho}_i(\cdot, \boldsymbol{\pi}_{-i}(\cdot; \boldsymbol{\omega}); \boldsymbol{\sigma}), \boldsymbol{\pi}_{-i}(\cdot, \boldsymbol{\omega}); \boldsymbol{h}_i^{\boldsymbol{\sigma}}) - \widehat{u}_i(\boldsymbol{\pi}(\cdot; \boldsymbol{\omega}); \boldsymbol{h})$, while the **cumulative regret gradient estimator** is given by $\widehat{G}(\boldsymbol{\omega}, \boldsymbol{\sigma}; \boldsymbol{h}, \boldsymbol{h}^{\boldsymbol{\sigma}}) \doteq (\nabla_{\boldsymbol{\omega}} \widehat{\Psi}(\boldsymbol{\omega}, \boldsymbol{\sigma}; \boldsymbol{h}, \boldsymbol{h}'), \nabla_{\boldsymbol{\sigma}} \widehat{\Psi}(\boldsymbol{\omega}, \boldsymbol{\sigma}; \boldsymbol{h}, \boldsymbol{h}^{\boldsymbol{\sigma}}))$.

Our main theorem requires one final definition, namely the **equilibrium distribution mismatch coefficient** $\left\| \frac{\partial \delta_{\mu}^{\boldsymbol{\pi}^*}}{\partial \mu} \right\|_{\infty}$, defined as the Radon-Nikodym derivative of the state-visitation distribution of the GNE $\boldsymbol{\pi}^*$ w.r.t. the initial state distribution μ. This coefficient, which measures the inherent difficulty of visiting states under the equilibrium policy $\boldsymbol{\pi}^*$—without knowing $\boldsymbol{\pi}^*$—is closely related to other distribution mismatch coefficients used in the analysis of policy gradient methods (Agarwal et al., 2020).

We now state our main theorem, namely that, under the assumptions outlined above, Algorithm 11 computes values for the policy parameters that nearly satisfy the necessary conditions for an MGPNE in polynomially many gradient steps, or equivalently, calls to the differentiable simulator.

Theorem 12.3.1.

Given a Markov pseudo-game $\mathcal{M}$, and a parameterization scheme $(\boldsymbol{\pi}, \boldsymbol{\rho}, \mathbb{R}^{\Omega}, \mathbb{R}^{\Sigma})$, suppose Assumption 12.2.1, 12.3.2, and 12.3.3 hold. For any $\delta > 0$, set $\varepsilon = \delta \|C_{br}(\cdot, \mu, \cdot)\|_{\infty}^{-1}$. If Algorithm 11 is run with inputs that satisfy, $\eta_{\boldsymbol{\omega}}, \eta_{\boldsymbol{\sigma}} \asymp \text{poly}(\varepsilon, \|\partial \delta_{\mu}^{\boldsymbol{\pi}^*} / \partial \mu\|_{\infty}, \frac{1}{1-\gamma}, \ell_{\nabla \Psi}^{-1}, \ell_{\Psi}^{-1})$, then for some $T \in \text{poly}\left(\varepsilon^{-1}, (1-\gamma)^{-1}, \|\partial \delta_{\mu}^{\boldsymbol{\pi}^*} / \partial \mu\|_{\infty}, \ell_{\nabla \Psi}, \ell_{\Psi}, \text{diam}(\mathbb{R}^{\Omega} \times \mathbb{R}^{\Sigma}), \eta_{\boldsymbol{\omega}}^{-1}\right)$, there exists $\boldsymbol{\omega}_{\text{best}}^{(T)} = \boldsymbol{\omega}^{(k)}$ with $k \leq T$ that is a $(\varepsilon, \varepsilon/2\ell_{\Psi})$-stationary point of the exploitability, i.e., there exists $\boldsymbol{\omega}^* \in \mathbb{R}^{\Omega}$ s.t. $\|\boldsymbol{\omega}_{\text{best}}^{(T)} - \boldsymbol{\omega}^*\| \leq \varepsilon/2\ell_{\Psi}$ and $\min_{\boldsymbol{h} \in \mathcal{D}\varphi(\boldsymbol{\omega}^*)} \|\boldsymbol{h}\| \leq \varepsilon$.

Further, for any arbitrary state distribution $\upsilon \in \Delta(\mathcal{S})$, if $\phi(\upsilon, \cdot)$ is differentiable at $\boldsymbol{\omega}^*$, $\|\nabla_{\boldsymbol{\omega}} \varphi(\upsilon, \boldsymbol{\omega}^*)\| \leq \delta$, i.e., $\boldsymbol{\omega}_{\text{best}}^{(T)}$ is a (ε, δ)-stationary point for the expected exploitability $\phi(\upsilon, \cdot)$.

Chapter 13

Radner Economies

13.1 Background

Having developed a mathematical formalism for Markov pseudo-games, along with a proof of existence of GMPE as well as an algorithm that computes them, we now move on to our main agenda, namely modeling incomplete stochastic economies in this formalism. We establish the first proof, to our knowledge, of the existence of recursive Radner equilibria in Radner economies, and we provide a polynomial-time algorithm for approximating them.

13.1.1 Static Exchange Economies

A **static exchange economy** (or **market**[1]) $(n, m, d, \mathcal{X}, \mathcal{E}, \mathcal{T}, \boldsymbol{u}, \boldsymbol{E}, \boldsymbol{\Theta})$, abbreviated by $(\boldsymbol{E}, \boldsymbol{\Theta})$ when clear from context, comprises a finite set of $n \in \mathbb{N}_+$ **consumers** and $m \in \mathbb{N}_+$ **commodities**. Each consumer $i \in [n]$ arrives at the market with an **endowment** of commodities represented as vector $\boldsymbol{e}_i = (e_{i1}, \dots, e_{im}) \in \mathcal{E}_i$, where $\mathcal{E}_i \subset \mathbb{R}^m$ is called the **endowment space**.[2] Any consumer i can sell its endowment $\boldsymbol{e}_i \in \mathcal{E}_i$ at **prices** $\boldsymbol{p} \in \mathcal{P}$, where $p_j \geq 0$ represents the value (resp. cost) of selling (resp. buying) a unit of commodity $j \in [m]$, to

[1] Although a static exchange "market" is an economy, we prefer the term "market" for the static components of a Radner economy, a dynamic exchange economy in which each time-period comprises one static market among many.

[2] Commodities are assumed to include labor services. Further, for any consumer i and endowment $\boldsymbol{e}_i \in \mathcal{E}_i$, $e_{ij} \geq 0$ denotes the quantity of commodity j in consumer i's possession, while $e_{ij} < 0$ denotes consumer i's debt, in terms of commodity j.

purchase a consumption $\boldsymbol{x}_i \in \mathcal{X}_i$ of commodities in its **consumption space** $\mathcal{X}_i \subseteq \mathbb{R}^m$.[3] Every consumer is constrained to buy a consumption with a cost weakly less than the value of its endowment, i.e., consumer i's **budget set**—the set of consumptions i can afford with its endowment $\boldsymbol{e}_i \in \mathcal{E}_i$ at prices $\boldsymbol{p} \in \mathcal{P}$—is determined by its **budget correspondence** $\mathcal{B}_i(\boldsymbol{e}_i, \boldsymbol{p}) \doteq \{\boldsymbol{x}_i \in \mathcal{X}_i \mid \boldsymbol{x}_i \cdot \boldsymbol{p} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}\}$.

Each consumer's consumption preferences are determined by its type-dependent preference relation $\succeq_{i,\boldsymbol{\theta}_i}$ on $\mathcal{X}_i$, represented by a type-dependent **utility function** $\boldsymbol{x}_i \mapsto u_i(\boldsymbol{x}_i; \boldsymbol{\theta}_i)$, for **type** $\boldsymbol{\theta}_i \in \mathcal{T}_i$ that characterizes consumer i's preferences within the **type space** $\mathcal{T}_i \subset \mathbb{R}^d$ of possible preferences.[4] The goal of each consumer i is thus to buy a consumption $\boldsymbol{x}_i \in \mathcal{B}_i(\boldsymbol{e}_i, \boldsymbol{p})$ that maximizes its utility function $\boldsymbol{x}_i \mapsto u_i(\boldsymbol{x}_i; \boldsymbol{\theta}_i)$ over its budget set $\mathcal{B}_i(\boldsymbol{e}_i, \boldsymbol{p})$.

We denote any **endowment profile** (resp. **type profile** and **consumption profile**) as $\boldsymbol{E} \doteq (\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n)^T \in \mathcal{E}$ (resp. $\boldsymbol{\Theta} \doteq (\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_n) \in \mathcal{T}$ and $\boldsymbol{X} \doteq (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)^T \in \mathcal{X}$). The **aggregate demand** (resp. **aggregate supply**) of a **consumption profile** $\boldsymbol{X} \in \mathcal{X}$ (resp. **an endowment profile** $\boldsymbol{E} \in \mathcal{E}$) is defined as the sum of consumptions (resp. endowments) across all consumers, i.e., $\sum_{i \in [n]} \boldsymbol{x}_i$ (resp. $\sum_{i \in [n]} \boldsymbol{e}_i$).

13.1.2 Radner Economies

A(n infinite horizon) Radner economy (Radner, 1972) $\mathcal{I} \doteq (n, m, l, d, \mathcal{S}, \mathcal{X}, \mathcal{E}, \mathcal{T}, \boldsymbol{u}, \gamma, \rho, \mathcal{R}, \mu)$, comprises $n \in \mathbb{N}$ consumers who, over an infinite discrete time horizon $t = 0, 1, 2, \ldots$, repeatedly encounter the opportunity to buy a consumption of $m \in \mathbb{N}$ commodities and a portfolio of $l \in \mathbb{N}$ assets, with their collective decisions leading them through a **state space** $\mathcal{S} \doteq \mathcal{O} \times (\mathcal{E} \times \mathcal{T})$. This state space comprises

[3]We note that, for any labor service j, consumer i's consumption x_{ij} is negative and restricted by its consumption space to be lower bounded by the negative of i's endowment, i.e., $x_{ij} \in [-e_{ij}, 0]$. This modeling choice allows us to model a consumer's preferences over the labor services she can provide. More generally, the consumption space models the constraints imposed on consumption by the "physical properties" of the world. That is, it rules out impossible combinations of commodities, such as strictly positive quantities of a commodity that is not available in the region where a consumer resides, or a supply of labor that amounts to more than 24 labor hours in a given day.

[4]In the sequel, we will be assuming, for any consumer i with any type $\boldsymbol{\theta}_i \in \mathcal{T}_i$, the type-dependent utility function $\boldsymbol{x}_i \mapsto u_i(\boldsymbol{x}_i; \boldsymbol{\theta}_i)$ is continuous, which implies that it can represent any type-dependent preference relation $\succeq_{i,\boldsymbol{\theta}_i}$ on $\mathbb{R}^m$ that is complete, transitive, and continuous (Debreu et al., 1954).

a **world state space** $\mathcal{O}$ and a **spot market space** $\mathcal{E} \times \mathcal{T}$. The spot market space is a collection of **spot markets**, each one a static exchange market $(\boldsymbol{E}, \boldsymbol{\Theta}) \in \mathcal{E} \times \mathcal{T} \subseteq \mathbb{R}^m \times \mathbb{R}^d$.

Each **asset** $k \in [l]$ is a **generalized Arrow security**, i.e., a divisible contract that transfers to its owner a quantity of the jth commodity at any world state $o \in \mathcal{O}$ determined by a matrix of asset returns $\boldsymbol{R}_o \doteq \left(\boldsymbol{r}_{o1}, \ldots, \boldsymbol{r}_{ol}\right)^T \in \mathbb{R}^{l \times m}$ s.t. $r_{okj} \in \mathbb{R}$ denotes the quantity of commodity j transferred at world state o for one unit of asset k. The collection of asset returns across all world states is given by $\mathcal{R} \doteq \{\boldsymbol{R}_o\}_{o \in \mathcal{O}}$. At any time step $t = 0, 1, 2, \ldots$, a consumer $i \in [n]$ can invest in an **asset portfolio** $\boldsymbol{y}_i \in \mathcal{Y}_i$ from a **space of asset portfolios (or investments)** $\mathcal{Y}_i \subset \mathbb{R}^l$ that define the **asset market**, where $y_{ik} \geq 0$ denotes the units of asset k bought (long) by consumer i, while $y_{ik} < 0$ denotes units that are sold (short). Assets are assumed to be **short-lived** (Magill and Quinzii, 1994), meaning that any asset purchased at time t pays its dividends in the subsequent time period $t + 1$, and then expires.[5] Assets allow consumers to insure themselves against future realizations of the spot market (i.e., types and endowments), by allowing it to transfer wealth across world states.

The economy starts at time period $t = 0$ in an **initial state** $S^{(0)} \sim \mu$ determined by an initial state distribution $\mu \in \Delta(\mathcal{S})$. At each time step $t = 0, 1, 2, \ldots$, the state of the economy is $\boldsymbol{s}^{(t)} \doteq (o^{(t)}, \boldsymbol{E}^{(t)}, \boldsymbol{\Theta}^{(t)}) \in \mathcal{S}$. Each consumer $i \in [n]$, observes the world state $o^{(t)} \in \mathcal{O}$, and participates in a spot market $(\boldsymbol{E}^{(t)}, \boldsymbol{\Theta}^{(t)})$, where it purchases **a consumption** $\boldsymbol{x}_i^{(t)} \in \mathcal{X}_i$ at **commodity prices** $\boldsymbol{p}^{(t)} \in \Delta_m$, and an **asset market** where it invests in an **asset portfolio** $\boldsymbol{y}_i^{(t)} \in \mathcal{Y}_i$ at **assets prices** $\boldsymbol{q}^{(t)} \in \mathbb{R}^l$.[6] Every consumer is constrained to buy a consumption $\boldsymbol{x}_i^{(t)} \in \mathcal{X}_i$ and invest in an asset portfolio $\boldsymbol{y}_i^{(t)} \in \mathcal{Y}_i$ with a total cost weakly less than the value of its current endowment $\boldsymbol{e}_i^{(t)} \in \mathcal{E}_i$. Formally, the set of consumptions and

[5]While for ease of exposition we assume that assets are short-lived, our results generalize to infinitely-lived generalized Arrow securities (Huang and Werner, 2004) (i.e., securities that never expire, so yield returns and can be resold in every subsequent time period following their purchase) with appropriate modifications to the definitions of the budget constraints and Walras' law. In contrast, our results do *not* immediately generalize to k-**period-living generalized Arrow securities** (i.e., securities that yield returns and can be resold in the k subsequent time periods following their purchase, until their expiration), as such securities introduce non-stationarities into the economy. To accommodate such securities would require that we generalize our Markov game model and methods to accommodate policies that depend on histories of length k.

[6]In general, asset prices can be negative. This modeling assumption is in line with the real world: e.g., it is common for energy futures to see negative prices because of costs associated with overproduction and limited storage capacity (Sheppard et al., 2020).

investment portfolios that a consumer i can afford with its current endowment $\boldsymbol{e}_i^{(t)} \in \mathcal{E}_i$ at current commodity prices $\boldsymbol{p}^{(t)} \in \mathcal{P}$ and current asset prices $\boldsymbol{q}^{(t)} \in \mathbb{R}^l$, i.e., its **budget set** $\mathcal{B}_i(\boldsymbol{e}_i^{(t)}, \boldsymbol{p}^{(t)}, \boldsymbol{q}^{(t)})$, is determined by its **budget correspondence**

$$\mathcal{B}_i(\boldsymbol{e}_i, \boldsymbol{p}, \boldsymbol{q}) \doteq \{(\boldsymbol{x}_i, \boldsymbol{y}_i) \in \mathcal{X}_i \times \mathcal{Y}_i \mid \boldsymbol{x}_i \cdot \boldsymbol{p} + \boldsymbol{y}_i \cdot \boldsymbol{q} \leq \boldsymbol{e}_i \cdot \boldsymbol{p}\}.$$

After the consumers make their consumption and investment decisions, they each receive **reward** $u_i(\boldsymbol{x}_i^{(t)}; \boldsymbol{\theta}_i^{(t)})$ as a function of their consumption and type, and then the economy either collapses with probability $1 - \gamma$, or survives with probability γ, where $\gamma \in (0, 1)$ is called the **discount rate**.[7] If the economy survives to see another day, then a new state is realized, namely $(O', E', \Theta') \sim \rho(\cdot \mid \boldsymbol{s}^{(t)}, \boldsymbol{Y}^{(t)})$, according to a **transition probability function** $\rho : \mathcal{S} \times \mathcal{S} \times \mathcal{Y} \to [0, 1]$ that depends on the consumers' investment portfolio profile $\boldsymbol{Y}^{(t)} \doteq (\boldsymbol{y}_1^{(t)}, \ldots, \boldsymbol{y}_n^{(t)})^T \in \mathcal{Y}$, after which the economy transitions to a new state $S^{(t+1)} \doteq (O', E' + \boldsymbol{Y}^{(t)} \boldsymbol{R}_{O'}, \Theta')$, where the consumers' new endowments depends on their returns $\boldsymbol{Y}^{(t)} \boldsymbol{R}_{O'} \in \mathbb{R}^{n \times m}$ on their investments.

Remark 13.1.1.

If only one commodity is delivered in exchange for assets, i.e., for all world states $o \in \mathcal{O}$, $\mathrm{Rank}(\boldsymbol{R}_o) \leq 1$, then the generalized Arrow securities are **numéraire generalized Arrow securities**, and the assets are called **financial assets**.[8] A Radner economy is **world-state-contingent** iff the cardinality of the world state space is weakly greater than that of the spot market space, i.e., $|\mathcal{O}| \geq |\mathcal{E} \times \mathcal{T}|$. Intuitively, when this condition holds, there exists a surjection from world states to spot market states, which implies that spot market states are implicit in world states, so that the spot market states can be dropped from the state space, i.e., $\mathcal{S} \doteq \mathcal{O}$. A Radner economy has **complete asset markets** if it is world-state-contingent,

[7]While for ease of exposition we assume a single discount factor for all consumers, our results extend to a setting in which each consumer $i \in [n]$ has a potentially unique discount factor $\gamma_i \in (0, 1)$ by incorporating the discount rates into the consumers' payoffs in the Markov pseudo-game defined in Section 14.1.2, rather than the history distribution.

[8]Recall that the numéraire is a fixed commodity that is used to standardize the value of other commodities, while a numéraire generalized Arrow security is a generalized Arrow security that delivers its returns in terms of the numérarire. If the assets deliver exactly one commodity, i.e., $\mathrm{Rank}(\boldsymbol{R}_o) = 1$ at all world states o, we take that commodity to be the numéraire for the corresponding spot markets. On the other hand, if the assets deliver no commodity, i.e., $\mathrm{Rank}(\boldsymbol{R}_o) = 0$ at world state o, then we can take any arbitrary commodity to be the numéraire, in which case, the assets vacuously "deliver" zero units of the numérarire, and no units of any other commodities either.

and assets can deliver some commodity at all world states, i.e., for all world states $o \in \mathcal{O}$, $\mathrm{Rank}(\boldsymbol{R}_o) \geq 1$. Otherwise, it has **incomplete asset markets**. Colloquially, we call an infinite horizon exchange economy with (in)complete asset markets an **(in)complete exchange economy**. Intuitively, in complete exchange economies, consumers can insure themselves against all future realizations of the spot market—uncertainty regarding their endowments and types—since a complete exchange economy is world-state contingent. Further, when there is only a single commodity, s.t. $m = 1$, and only one financial asset which is a risk-free bond s.t. $l = 1$, and the return matrix for all world states $o \in \mathcal{O}$ (now a scalar since there is only one commodity and one financial asset) is given by $r_o \doteq \alpha$, for some $\alpha \in \mathbb{R}$, we obtain the standard incomplete market model (Blackwell, 1965; Lucas Jr and Prescott, 1971).

A **history** $\boldsymbol{h} \in \mathcal{H}^\tau \doteq (\mathcal{S} \times \mathcal{X} \times \mathcal{Y} \times \mathcal{P} \times \mathbb{R}^l)^\tau \times \mathcal{S}$ is a sequence $\boldsymbol{h} = ((\boldsymbol{s}^{(t)}, \boldsymbol{X}^{(t)}, \boldsymbol{Y}^{(t)}, \boldsymbol{p}^{(t)}, \boldsymbol{q}^{(t)})_{t=0}^{\tau-1}, \boldsymbol{s}^{(\tau)})$ of tuples comprising states, consumption profiles, investment profiles, commodity price, and asset prices s.t. a history of length 0 corresponds only to the initial state of the economy. For any history $\boldsymbol{h} \in \mathcal{H}^\tau$, we denote by $\boldsymbol{h}_{:p}$ the first $p \in [0:\tau]$ steps of $\boldsymbol{h}$, i.e., $\boldsymbol{h}_{:p} \doteq ((\boldsymbol{s}^{(t)}, \boldsymbol{X}^{(t)}, \boldsymbol{Y}^{(t)}, \boldsymbol{p}^{(t)}, \boldsymbol{q}^{(t)})_{t=0}^{p-1}, \boldsymbol{s}^{(p)})$. Overloading notation, we define the **history space** $\mathcal{H} \doteq \bigcup_{\tau=0}^{\infty} \mathcal{H}^\tau$, and then **consumption**, **investment**, **commodity price** and **asset price policies** as mappings $\boldsymbol{x}_i : \mathcal{H} \to \mathcal{X}_i$, $\boldsymbol{y}_i : \mathcal{H} \to \mathcal{Y}_i$, $\boldsymbol{p} : \mathcal{H} \to \Delta_m$, and $\boldsymbol{q} : \mathcal{H} \to \mathbb{R}^l$ from histories to consumptions, investments, commodity prices, and asset prices, respectively, s.t. $(\boldsymbol{x}_i, \boldsymbol{y}_i)(\boldsymbol{h})$ is the consumption-investment decision of consumer $i \in [n]$, and $(\boldsymbol{p}, \boldsymbol{q})(\boldsymbol{h})$ are commodity and asset prices, both at history $\boldsymbol{h} \in \mathcal{H}$. A **consumption policy profile** (resp. **investment policy profile**) $\boldsymbol{X}(\boldsymbol{h}) \doteq (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)(\boldsymbol{h})^T$ (resp. $\boldsymbol{Y}(\boldsymbol{h}) \doteq (\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n)(\boldsymbol{h})^T$) is a collection of consumption (resp. investment) policies for all consumers. A consumption policy $\boldsymbol{x}_i : \mathcal{S} \to \mathcal{X}_i$ is **Markov** if it depends only on the last state of the history, i.e., $\boldsymbol{x}_i(\boldsymbol{h}) = \boldsymbol{x}_i(\boldsymbol{s}^{(\tau)})$, for all histories $\boldsymbol{h} \in \mathcal{H}^\tau$ of all lengths $\tau \in \mathbb{N}$. An analogous definition extends to investment, commodity price, and asset price policies.

Given $\boldsymbol{\pi} \doteq (\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{p}, \boldsymbol{q})$ and a history $\boldsymbol{h} \in \mathcal{H}^\tau$, we define the **discounted history distribution** assuming initial state distribution μ as

$$\nu_\mu^{\boldsymbol{\pi},\tau}(\boldsymbol{h}) = \mu(\boldsymbol{s}^{(0)}) \prod_{t=0}^{\tau-1} \gamma^t \rho(o^{(t+1)}, \boldsymbol{E}^{(t+1)} + \boldsymbol{Y}^{(t)} \boldsymbol{R}_{o^{(t+1)}}, \boldsymbol{\Theta}^{(t+1)} \mid \boldsymbol{s}^{(t)}, \boldsymbol{Y}^{(t)}) \mathbb{1}_{\{\boldsymbol{Y}(\boldsymbol{h}_{:t})\}}(\boldsymbol{Y}^{(t)}).$$

Overloading notation, we define the set of all realizable trajectories $\mathcal{H}^{\boldsymbol{\pi}}$ of length τ under policy profile $\boldsymbol{\pi}$ as $\mathcal{H}^{\boldsymbol{\pi}} \doteq \operatorname{supp}(\nu_\mu^{\boldsymbol{\pi},\tau})$, i.e., the set of all histories that occur with non-zero probability, and we let $H = \left((S^{(t)}, A^{(t)})_{t=0}^{\tau-1}, S^{(\tau)}\right)$ be any randomly sampled history from $\nu_\mu^{\boldsymbol{\pi},\tau}$. Finally, we abbreviate $\nu_\mu^{\boldsymbol{\pi}} \doteq \nu_\mu^{\boldsymbol{\pi},\infty}$.

13.1.3 Solution Concepts and Existence

An **outcome** $(\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{p}, \boldsymbol{q}) : \mathcal{H} \to \mathcal{X} \times \mathcal{Y} \times \Delta_m \times \mathbb{R}^l$ of a Radner economy is a tuple consisting of a commodity prices policy, an asset prices policy, a consumption policy profile, and an investment policy profile.[9]

An outcome is **Markov** if all its constituent policies are Markov: i.e., if it depends only on the last state of the history, i.e., $(\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{p}, \boldsymbol{q})(\boldsymbol{h}) = (\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{p}, \boldsymbol{q})(\boldsymbol{s}^{(\tau)})$, for all histories $\boldsymbol{h} \in \mathcal{H}^\tau$ of all lengths $\tau \in \mathbb{N}$.

We now introduce a number of properties of Radner economies outcomes, which we use to define our solution concepts. While these properties are defined broadly for (in general, history-dependent) outcomes, they also apply in the special case of Markov outcomes.

Given a consumption and investment profile $(\boldsymbol{X}, \boldsymbol{Y})$, the **consumption state-value function** $v_i^{(\boldsymbol{X},\boldsymbol{Y},\boldsymbol{p},\boldsymbol{q})} : \mathcal{S} \to \mathbb{R}$ is defined as:

$$v_i^{(\boldsymbol{X},\boldsymbol{Y},\boldsymbol{p},\boldsymbol{q})}(\boldsymbol{s}) \doteq \mathop{\mathbb{E}}_{H \sim \nu_\mu^{(\boldsymbol{X},\boldsymbol{Y},\boldsymbol{p},\boldsymbol{q})}} \left[\sum_{t=0}^{\infty} \gamma^t u_i \left(\boldsymbol{x}_i(H_{:t}); \Theta^{(t)} \right) \mid S^{(0)} = \boldsymbol{s} \right].$$

Definition 13.1.1 [Optimal Outcome]**.**

An outcome $(\boldsymbol{X}^*, \boldsymbol{Y}^*, \boldsymbol{p}^*, \boldsymbol{q}^*)$ is **optimal** for i if i's **expected cumulative utility** $u_i(\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{p}, \boldsymbol{q}) \doteq \mathbb{E}_{\boldsymbol{s} \sim \mu} \left[v_i^{(\boldsymbol{X},\boldsymbol{Y},\boldsymbol{p},\boldsymbol{q})}(\boldsymbol{s}) \right]$ is maximized over all affordable consumption and

[9] Instead of expressing this tuple as $\mathcal{X}^{\mathcal{H}} \times \mathcal{Y}^{\mathcal{H}} \times {\Delta_m}^{\mathcal{H}} \times {\mathbb{R}^l}^{\mathcal{H}}$, we sometimes write $(\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{p}, \boldsymbol{q}) : \mathcal{H} \to \mathcal{X} \times \mathcal{Y} \times \Delta_m \times \mathbb{R}^l$.

investment policies, i.e.,

$$(\boldsymbol{x}_i^*, \boldsymbol{y}_i^*) \in \mathop{\arg\max}_{\substack{(\boldsymbol{x}_i, \boldsymbol{y}_i): \mathcal{H} \to \mathcal{X}_i \times \mathcal{Y}_i, \forall t \in \mathbb{N}, \boldsymbol{h} \in \mathcal{H}^t \\ (\boldsymbol{x}_i, \boldsymbol{y}_i)(\boldsymbol{h}_{:t}) \in \mathcal{B}_i(\boldsymbol{e}_i^{(t)}, \boldsymbol{p}^*(\boldsymbol{h}_{:t}), \boldsymbol{q}^*(\boldsymbol{h}_{:t}))}} u_i(\boldsymbol{x}_i, \boldsymbol{x}_{-i}^*, \boldsymbol{y}_i, \boldsymbol{y}_{-i}^*, \boldsymbol{p}^*, \boldsymbol{q}^*) \ . \tag{13.1}$$

A Markov outcome $(\boldsymbol{X}^*, \boldsymbol{Y}^*, \boldsymbol{p}^*, \boldsymbol{q}^*)$ is **Markov perfect** for i if i maximizes its consumption state-value function over all affordable consumption and investment policies, i.e.,

$$(\boldsymbol{x}_i^*, \boldsymbol{y}_i^*) \in \mathop{\arg\max}_{\substack{(\boldsymbol{x}_i, \boldsymbol{y}_i): \mathcal{S} \to \mathcal{X}_i \times \mathcal{Y}_i : \forall \boldsymbol{s} \in \mathcal{S}, \\ (\boldsymbol{x}_i, \boldsymbol{y}_i)(\boldsymbol{s}) \in \mathcal{B}_i(\boldsymbol{e}_i, \boldsymbol{p}^*(\boldsymbol{s}), \boldsymbol{q}^*(\boldsymbol{s}))}} \left\{ v_i^{(\boldsymbol{x}_i, \boldsymbol{x}_{-i}^*, \boldsymbol{y}_i, \boldsymbol{y}_{-i}^*, \boldsymbol{p}^*, \boldsymbol{q}^*)}(\boldsymbol{s}) \right\} \ . \tag{13.2}$$

Definition 13.1.2 [Feasible Outcomes]**.**

A consumption policy $\boldsymbol{X}$ is said to be **feasible** iff for all time horizons $\tau \in \mathbb{N}$ and histories $\boldsymbol{h} \in \mathcal{H}^\tau$ of length τ,

$$\sum_{i \in [n]} \boldsymbol{x}_i(\boldsymbol{h}) - \sum_{i \in [n]} \boldsymbol{e}_i^{(\tau)} \leq \boldsymbol{0}_m,$$

where $\boldsymbol{e}_i^{(\tau)} \in \mathcal{E}_i$ is consumer i's endowment at the end of history $\boldsymbol{h}$, i.e., at state $\boldsymbol{s}^{(\tau)}$.

Similarly, an investment policy is **feasible** iff for all time horizons $\tau \in \mathbb{N}$ and histories $\boldsymbol{h} \in \mathcal{H}^\tau$ of length τ,

$$\sum_{i \in [n]} \boldsymbol{y}_i(\boldsymbol{h}) \leq \boldsymbol{0}_l.$$

If all the consumption and investment policies associated with an outcome are feasible, we will colloquially refer to the outcome as **feasible** as well.

Definition 13.1.3 [Walras' Law]**.**

An outcome $(\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{p}, \boldsymbol{q})$ is said to satisfy **Walras' law** iff for all time horizons $\tau \in \mathbb{N}$ and histories $\boldsymbol{h} \in \mathcal{H}^\tau$ of length τ,

$$\boldsymbol{p}(\boldsymbol{h}) \cdot \left(\sum_{i \in [n]} \boldsymbol{x}_i(\boldsymbol{h}) - \sum_{i \in [n]} \boldsymbol{e}_i^{(\tau)} \right) + \boldsymbol{q}(\boldsymbol{h}) \cdot \left(\sum_{i \in [n]} \boldsymbol{y}_i(\boldsymbol{h}) \right) = 0,$$

where, as above, $\boldsymbol{e}_i^{(\tau)} \in \mathcal{E}_i$ is consumer i's endowment at state $\boldsymbol{s}^{(\tau)}$.

The canonical solution concept for stochastic economies is the Radner equilibrium.

Definition 13.1.4 [Radner Equilibrium]**.**

A **Radner (or sequential competitive) equilibrium (RE)** (Radner, 1972) of a Radner economy $\mathcal{I}$ is an outcome $(\boldsymbol{X}^*, \boldsymbol{Y}^*, \boldsymbol{p}^*, \boldsymbol{q}^*)$ that is 1. optimal for all consumers, i.e., Equation (13.1) is satisfied, for all consumers $i \in [n]$; 2. feasible; and 3. satisfies Walras' law.

As a Radner equilibrium is in general infinite dimensional, we are often interested in a recursive Radner equilibrium which is a *Markov* outcome, i.e., one that depends only on the last state of the history rather than the entire history, and as such better behaved.

Definition 13.1.5 [Recursive Radner Equilibrium]**.**

A **recursive Radner (or Walrasian or competitive) equilibrium (RRE)** (Mehra and Prescott, 1977; Prescott and Mehra, 1980) of a Radner economy $\mathcal{I}$ is a Markov outcome $(\boldsymbol{X}^*, \boldsymbol{Y}^*, \boldsymbol{p}^*, \boldsymbol{q}^*)$ that is 1. Markov perfect for all consumers, i.e., Equation (13.2) is satisfied, for all consumers $i \in [n]$; 2. feasible; and 3. satisfies Walras' law.

The following assumptions are standard in the equilibrium literature (see, for instance, Geanakoplos (1990)). We prove the existence of a recursive Radner equilibrium under these assumptions.

Assumption 13.1.1.

Given a Radner economy $\mathcal{I}$, assume for all $i \in [n]$,

1. $\mathcal{X}$, $\mathcal{Y}$, $\mathcal{E}$, are non-empty, closed, convex, with $\mathcal{E}$ additionally bounded;
2. $(\boldsymbol{\theta}_i, \boldsymbol{x}_i) \mapsto u_i(\boldsymbol{x}_i; \boldsymbol{\theta}_i)$ is continuous and concave, and $(\boldsymbol{s}, \boldsymbol{y}_i) \mapsto \rho(\boldsymbol{s}' \mid \boldsymbol{s}, \boldsymbol{y}_i, \boldsymbol{y}_{-i})$ is continuous and stochastically concave, for all $\boldsymbol{s}, \in \mathcal{S}$ and $\boldsymbol{y}_{-i} \in \mathcal{Y}_{-i}$;
3. for all $\boldsymbol{e}_i \in \mathcal{E}_i$, the correspondence

$$(\boldsymbol{p}, \boldsymbol{q}) \rightrightarrows \mathcal{B}_i(\boldsymbol{e}_i, \boldsymbol{p}, \boldsymbol{q}) \cap \{(\boldsymbol{x}_i, \boldsymbol{y}_i) \mid \sum_{i \in [n]} \boldsymbol{x}_i \leq \sum_{i \in [n]} \boldsymbol{e}_i, \sum_{i \in [n]} \boldsymbol{y}_i \leq \boldsymbol{0}_m, (\boldsymbol{X}, \boldsymbol{Y}) \in \mathcal{X} \times \mathcal{Y}\}$$

is continuous[10];

[10]One way to ensure that this condition holds is to assume that for all $\boldsymbol{s} = (o, \boldsymbol{E}, \boldsymbol{\Theta}) \in \mathcal{S}$, returns from assets are positive $\boldsymbol{R}_o \geq \boldsymbol{0}_{ml}$, and for all consumers $i \in [n]$, there exists $(\boldsymbol{x}_i, \boldsymbol{y}_i) \in \mathcal{X}_i \times \mathcal{Y}_i$, s.t. $\boldsymbol{x}_i < \boldsymbol{e}_i$, $\boldsymbol{y}_i < 0$.

4. $\mathcal{B}_i(\boldsymbol{e}_i, \boldsymbol{p}, \boldsymbol{q}) \cap \{(\boldsymbol{x}_i, \boldsymbol{y}_i) \mid \sum_{i\in[n]} \boldsymbol{x}_i \leq \sum_{i\in[n]} \boldsymbol{e}_i, \sum_{i\in[n]} \boldsymbol{y}_i \leq \boldsymbol{0}_m, (\boldsymbol{X}, \boldsymbol{Y}) \in \mathcal{X} \times \mathcal{Y}\}$ is non-empty, convex, and compact, for all $\boldsymbol{e}_i \in \mathcal{E}_i$, $\boldsymbol{p} \in \Delta_m$, and $\boldsymbol{q} \in \mathbb{R}^l$[11];
5. (no saturation) there exists an $\boldsymbol{x}_i^+ \in \mathcal{X}_i$ s.t. $u_i(\boldsymbol{x}_i^+; \boldsymbol{\theta}_i) > u_i(\boldsymbol{x}_i; \boldsymbol{\theta}_i)$, for all $\boldsymbol{x}_i \in \mathcal{X}_i$ and $\boldsymbol{\theta}_i \in \mathcal{T}_i$.

Next we associate an **Radner Markov pseudo-game** $\mathcal{M}$ with a given Radner economy $\mathcal{I}$.

Definition 13.1.6 [Radner Markov pseudo-game]**.**

Let $\mathcal{I}$ be a Radner economy. The corresponding **Radner Markov pseudo-game** $\mathcal{M} = (n+1, m+l, \mathcal{S}, \bigtimes_{i\in[n]}(\mathcal{X}_i \times \mathcal{Y}_i) \times (\mathcal{P} \times \mathcal{Q}), \mathcal{B}', r', \rho', \gamma', \mu')$ is defined as

- The $n+1$ players comprise n consumers, players $1, \ldots, n$, and one auctioneer, player $n+1$.
- The set of states $\mathcal{S} = \mathcal{O} \times \mathcal{E} \times \mathcal{T}$. At each state $\boldsymbol{s} = (o, \boldsymbol{E}, \boldsymbol{\Theta}) \in \mathcal{S}$,
  - each consumer $i \in [n]$ chooses an action $\boldsymbol{a}_i = (\boldsymbol{x}_i, \boldsymbol{y}_i) \in \mathcal{B}'_i\left(\boldsymbol{s}, \boldsymbol{a}_{-i}\right) \subseteq \mathcal{X}_i \times \mathcal{Y}_i$ from a set of feasible actions $\mathcal{B}'_i(\boldsymbol{s}, \boldsymbol{a}_{-i}) = \mathcal{B}_i(\boldsymbol{e}_i, \boldsymbol{a}_{n+1}) \cap \{(\boldsymbol{x}_i, \boldsymbol{y}_i) \mid \sum_{i\in[n]} \boldsymbol{x}_i \leq \sum_{i\in[n]} \boldsymbol{e}_i, \sum_{i\in[n]} \boldsymbol{y}_i \leq \boldsymbol{0}_m, (\boldsymbol{X}, \boldsymbol{Y}) \in \mathcal{X} \times \mathcal{Y}\}$ and receives reward $r'_i(\boldsymbol{s}, \boldsymbol{a}) \doteq u_i(\boldsymbol{x}_i; \boldsymbol{\theta}_i)$; and
  - the auctioneer $n+1$ chooses an action $\boldsymbol{a}_{n+1} = (\boldsymbol{p}, \boldsymbol{q}) \in \mathcal{B}'_{n+1}\left(\boldsymbol{s}, \boldsymbol{a}_{-(n+1)}\right) \doteq \mathcal{P} \times \mathcal{Q}$ where $\mathcal{P} \doteq \Delta_m$ and $\mathcal{Q} \subseteq [0, \max_{\boldsymbol{E}\in\mathcal{E}} \sum_{i\in[n]} \sum_{j\in[m]} e_{ij}]^l$, and and receives reward $r'_{n+1}(\boldsymbol{s}, \boldsymbol{a}) \doteq \boldsymbol{p} \cdot \left(\sum_{i\in[n]} \boldsymbol{x}_i - \sum_{i\in[n]} \boldsymbol{e}_i\right) + \boldsymbol{q} \cdot \left(\sum_{i\in[n]} \boldsymbol{y}_i\right)$.
- The transition probability function is defined as $\rho'(\boldsymbol{s}' \mid \boldsymbol{s}, \boldsymbol{a}) \doteq \rho(\boldsymbol{s}' \mid \boldsymbol{s}, \boldsymbol{Y})$.
- The discount rate $\gamma' = \gamma$ and the initial state distribution $\mu' = \mu$.

Our existence proof reformulates the set of recursive Radner equilibria of any Radner economy as the set of GMPE of the Radner Markov pseudo-game.

Theorem 13.1.1.

Consider a Radner economy $\mathcal{I}$. Under Assumption 13.1.1, the set of recursive Radner

[11]One way to ensure that this condition holds is to assume that for all $\boldsymbol{s} = (o, \boldsymbol{E}, \boldsymbol{\Theta}) \in \mathcal{S}$, returns from assets are positive, i.e., $\boldsymbol{R}_o \geq \boldsymbol{0}_{ml}$, and $\mathcal{X}, \mathcal{Y}$ are bounded from below.

equilibria of $\mathcal{I}$ is equal to the set of GMPE of the associated Radner Markov pseudo-game $\mathcal{M}$.

Corollary 13.1.1.

Under Assumption 13.1.1, the set of recursive Radner equilibria of a Radner economy is non-empty.

13.1.4 Equilibrium Computation

Since a recursive Radner equilibrium is infinite-dimensional when the state space is continuous, its computation is FNP-hard (Murty and Kabadi, 1987). As such, it is generally believed that the best we can hope to find in polynomial time is an outcome that satisfies the necessary conditions of a stationary point of a recursive Radner equilibrium. Since the set of recursive Radner equilibria of any Radner economy is equal to the set of GMPE of the associated Radner Markov pseudo-game (Theorem 13.1.1), running Algorithm 11 on this Radner Markov pseudo-game will allow us to compute a policy profile that satisfies the necessary conditions of a stationary point of an GMPE, and hence a recursive Radner equilibrium.

Combining Theorem 13.1.1 and Theorem 12.3.1, we thus obtain the following computational complexity guarantees for Algorithm 11, when run on the Radner Markov pseudo-game associated with a Radner economy.[12]

Theorem 13.1.2.

Consider a Radner economy $\mathcal{I}$ and the associated Radner Markov pseudo-games $\mathcal{M}$. Let $(\boldsymbol{\pi}, \boldsymbol{\rho}, \mathbb{R}^{\Omega}, \mathbb{R}^{\Sigma})$ be a parametrization scheme for $\mathcal{M}$ and suppose Assumptions 12.3.2, 12.3.3, and 13.1.1 hold. Then, the convergence results in Theorem 12.3.1 hold for $\mathcal{M}$.

[12]While for generality and ease of exposition we state Assumptions 12.3.2 and 12.3.3 for the Radner Markov pseudo-game $\mathcal{M}$, we note that when the Radner economy $\mathcal{I}$ satisfies Assumption 13.1.1, to ensure that the associated Radner Markov pseudo-game $\mathcal{M}$ satisfies Assumption 12.3.2 and 12.3.3, it suffices to assume that the parametric policy functions $(\boldsymbol{\pi}, \boldsymbol{\rho})$ are affine; the policy parameter spaces $(\mathbb{R}^{\Omega}, \mathbb{R}^{\Sigma})$ are non-empty, compact, and convex; for all players $i \in [n]$ and types $\boldsymbol{\theta}_i \in \mathcal{T}_i$, the utility function $\boldsymbol{x}_i \mapsto u_i(\boldsymbol{x}_i; \boldsymbol{\theta}_i)$ is twice continuously differentiable; and for all $\boldsymbol{s}, \boldsymbol{s}, \in \mathcal{S}$, the transition function $\boldsymbol{Y} \mapsto \rho(\boldsymbol{s}' \mid \boldsymbol{s}, \boldsymbol{Y})$ is twice continuously differentiable.

13.2 Experiments

Given a Radner economy $\mathcal{I}$, we associate with it an exchange economy Markov pseudo-game $\mathcal{M}$, and we then construct a neural network to solve $\mathcal{M}$. To do so, we assume a parametrization scheme $(\boldsymbol{\pi}, \boldsymbol{\rho}, \mathbb{R}^{\Omega}, \mathbb{R}^{\Sigma})$, where the parametric policy profiles $(\boldsymbol{\pi}, \boldsymbol{\rho})$ are represented by neural networks with $(\mathbb{R}^{\Omega}, \mathbb{R}^{\Sigma})$ as the corresponding network weights. Computing an RRE via Algorithm 11 can then be seen as the result of training a generative adversarial neural network (Goodfellow et al., 2014), where $\boldsymbol{\pi}$ (resp. $\boldsymbol{\rho}$) is the output of the generator (resp. adversarial) network. We call such a neural representation a **generative adversarial policy network (GAPNet)**.

Following this approach, we built GAPNets to approximate the RRE in two types of Radner economies: one with a deterministic transition probability function and another with a stochastic transition probability function. Within each type, we experimented with three randomly sampled economies, each with 10 consumers, 10 commodities, 1 asset, 5 world states, and characterized by a distinct class of reward functions, which impart different smoothness properties onto the state-value function: **linear**: $u_i(\boldsymbol{x}_i; \boldsymbol{\theta}_i) = \sum_{j \in [m]} \theta_{ij} x_{ij}$; **Cobb-Douglas**: $u_i(\boldsymbol{x}_i; \boldsymbol{\theta}_i) = \prod_{j \in [m]} {x_{ij}}^{\theta_{ij}}$; and **Leontief**: $u_i(\boldsymbol{x}_i; \boldsymbol{\theta}_i) = \min_{j \in [m]} \left\{ \frac{x_{ij}}{\theta_{ij}} \right\}$.[13]

We compare our results with a classic neural projection method (also known as deep equilibrium nets (Azinovic et al., 2022)), which macroeconomists and others use to solve stochastic economies. In this latter method, one seeks a policy profile that minimizes the norm of the system of first-order necessary and sufficient conditions that characterize RRE (see for instance, (Fernández-Villaverde, 2023)).[14] We use the same network architecture for both methods, and select hyperparameters through grid search. In all experiments, we evaluate the performance of the computed policy profiles using three metrics: total first-order violations, average Bellman errors,[15] and exploitability. For each metric, we

[13] Full details of our experimental setup appear in Section 14.2, including hyperparameter search values, final experimental configurations, and visualization code. See also our code repository: `https://github.com/Sadie-Zhao/Markov-Pseudo-Game-EC2025`.

[14] We describe the neural projection method in Section 14.2.1.

[15] The definitions of these two metrics can be found in Section 14.2.1.

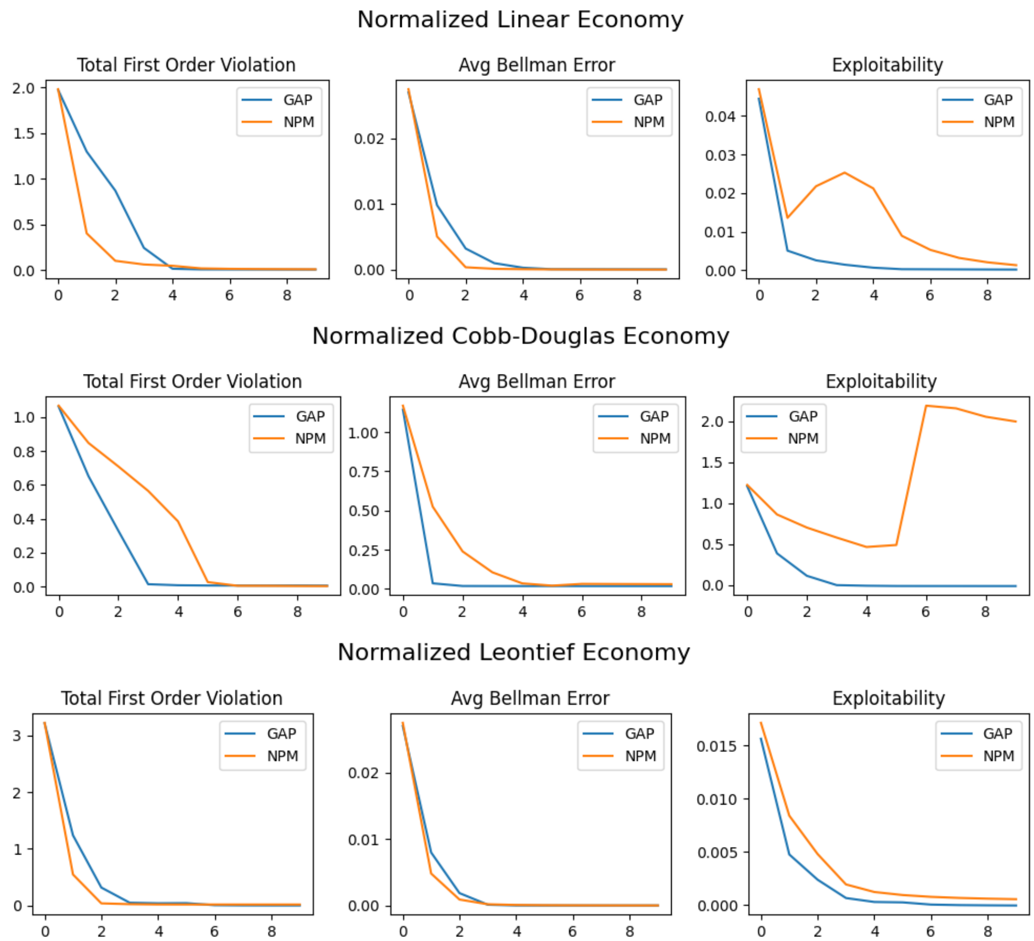

Figure 13.1: Normalized Metrics for Economies with Deterministic Transition Probability Function

randomly sample 50 policy profiles, record their corresponding values, and normalize the results by dividing it by the average.

Figure 13.1 depicts our results for economies with deterministic transition functions. Perhaps unsurprisingly, while GAPNets demonstrates a clear advantage in minimizing exploitability in all three economies, the neural projection method (NPM) slightly outperforms GAPNets in minimizing total first order violations and average Bellman error, the metrics they are specifically designed to minimize. Furthermore, in all three economies, exploitability is near 0, highlighting GAPNet's ability to approximate at least a Radner equilibrium. Figure 13.2 presents our results for economies with stochastic transition functions. These results indicate that stochasticity hinders NPM's ability to minimize the three metrics, even

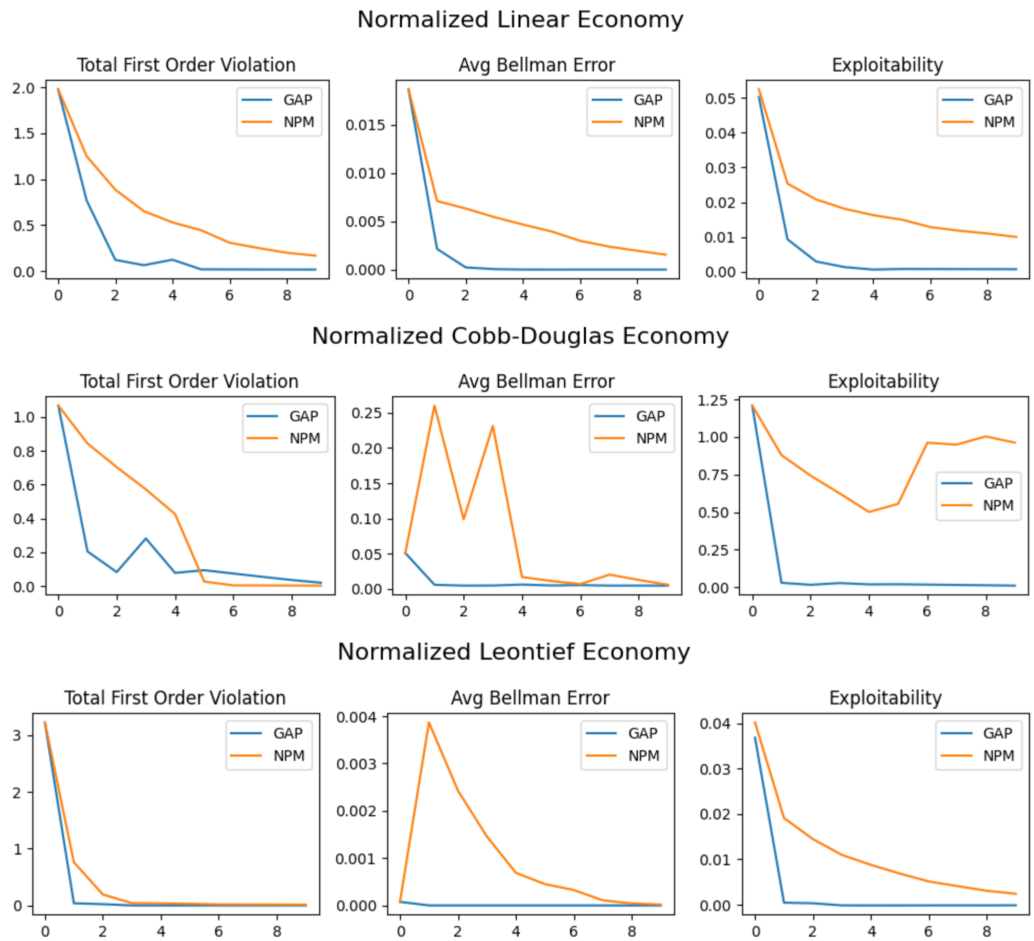

Figure 13.2: Normalized Metrics for Radner Economies with Stochastic Transition Probability Function

the method is explictly designed to minimize two of them. In contrast, GAPNet successfully minimizes all three metrics across all economies.

Chapter 14

Appendix for Part III

14.1 Omitted Results and Proofs

14.1.1 Omitted Results and Proofs from Chapter 12

Theorem 12.2.1.

Let $\mathcal{M}$ be a Markov pseudo-game for which Assumption 12.2.1 holds, and let $\mathcal{P}^{\mathrm{sub}} \subseteq \mathcal{P}^{\mathrm{markov}}$ be a subspace of Markov policy profiles that satisfies Assumption 12.2.2. Then, there exists a policy $\boldsymbol{\pi}^* \in \mathcal{P}^{\mathrm{sub}}$ such that $\boldsymbol{\pi}^*$ is an GMPE of $\mathcal{M}$.

Proof

First, by Part 3 of Assumption 12.2.1, we know that for any $i \in [n]$, $\mathcal{F}_i^{\mathrm{sub}}(\boldsymbol{\pi}_{-i})$ is non-empty, convex, and compact, for all $\boldsymbol{\pi}_{-i} \in \mathcal{P}_{-i}$. Moreover, 2 of Assumption 12.2.1, $\mathcal{F}^{\mathrm{sub}}$ is upper-hemicontinuous. Therefore, by the Kakutani-Glicksberg fixed-point theorem (see, Theorem 2.4.1—(Glicksberg, 1952)), the set $\mathcal{F}^{\mathrm{sub}} \doteq \{\boldsymbol{\pi} \in \mathcal{P}^{\mathrm{sub}} \mid \boldsymbol{\pi} \in \mathcal{F}^{\mathrm{sub}}(\boldsymbol{\pi})\}$ is non-empty.

For any player $i \in [n]$ and state $\boldsymbol{s} \in \mathcal{S}$, we define the **individual state best-response correspondence** $\Phi_i^{\boldsymbol{s}} : \mathcal{P}^{\mathrm{sub}} \rightrightarrows \mathcal{A}_i$ by

$$\Phi_i^{\boldsymbol{s}}(\boldsymbol{\pi}) \doteq \operatorname*{arg\,max}_{\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{s}, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))} r_i(\boldsymbol{s}, \boldsymbol{a}_i, \boldsymbol{\pi}_{-i}(\boldsymbol{s})) + \mathop{\mathbb{E}}_{S' \sim \rho(\cdot \mid \boldsymbol{s}, \boldsymbol{a}_i, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))} [\gamma v_i^{\boldsymbol{\pi}}(S')] \tag{14.1}$$

$$= \operatorname*{arg\,max}_{\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{s}, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))} q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{a}_i, \boldsymbol{\pi}_{-i}(\boldsymbol{s})) \tag{14.2}$$

Then, for any player $i \in [n]$, we define the **restricted individual best-response correspondence** $\Phi_i : \mathcal{P}^{\text{sub}} \rightrightarrows \mathcal{P}_i^{\text{sub}}$ as the Cartesian product of individual state best-response correspondences across the states restricted to $\mathcal{P}^{\text{sub}}$:

$$\Phi_i(\boldsymbol{\pi}) = \left(\bigtimes_{\boldsymbol{s} \in \mathcal{S}} \Phi_i^{\boldsymbol{s}}(\boldsymbol{\pi}) \right) \bigcap \mathcal{P}_i^{\text{sub}} \tag{14.3}$$

$$= \{\boldsymbol{\pi}_i \in \mathcal{P}_i^{\text{sub}} \mid \boldsymbol{\pi}_i(\boldsymbol{s}) \in \Phi_i^{\boldsymbol{s}}(\boldsymbol{\pi}), \forall\, \boldsymbol{s} \in \mathcal{S}\} \tag{14.4}$$

Finally, we can define the **multi-player best-response correspondence** $\Phi : \mathcal{P}^{\text{sub}} \rightrightarrows \mathcal{P}^{\text{sub}}$ as the Cartesian product of the individual correspondences, i.e., $\Phi(\boldsymbol{\pi}) \doteq \bigtimes_{i \in [n]} \Phi_i(\boldsymbol{\pi})$.

To show the existence of

MPGNE, we first want to show that there exists a fixed point $\boldsymbol{\pi}^* \in \mathcal{P}^{\text{sub}}$ such that $\boldsymbol{\pi}^* \in \Phi(\boldsymbol{\pi}^*)$. To this end, we need to show that 1. for any $\boldsymbol{\pi} \in \mathcal{P}^{\text{sub}}$, $\Phi(\boldsymbol{\pi})$ is non-empty, compact, and convex; 2. Φ is upper hemicontinuous.

Take any $\boldsymbol{\pi} \in \mathcal{P}^{\text{sub}}$. Fix $i \in [n], \boldsymbol{s} \in \mathcal{S}$, we know that $\boldsymbol{a}_i \mapsto q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{a}_i, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))$ is concave over $\mathcal{X}_i(\boldsymbol{s}, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))$, and $\mathcal{X}_i(\boldsymbol{s}, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))$ is non-empty, convex, and compact by Assumption 12.2.1, then by Proposition 4.1 of Fiacco and Kyparisis (1986), $\Phi_i^{\boldsymbol{s}}(\boldsymbol{\pi})$ is non-empty, compact, and convex.

Now, notice $\bigtimes_{\boldsymbol{s} \in \mathcal{S}} \Phi_i^{\boldsymbol{s}}(\boldsymbol{\pi})$ is compact and convex as it is a Cartesian product of compact, convex sets. Thus, as $\mathcal{P}^{\text{sub}}$ is also compact and convex by Assumption 12.2.2, we know that $\Phi_i(\boldsymbol{\pi}) = \left(\bigtimes_{\boldsymbol{s} \in \mathcal{S}} \Phi_i^{\boldsymbol{s}}(\boldsymbol{\pi}) \right) \bigcap \mathcal{P}_i^{\text{sub}}$ is compact and convex. By the assumption of *closure under policy improvement* under Assumption 12.2.2, we know that since $\boldsymbol{\pi} \in \mathcal{P}^{\text{sub}}$, there exists $\boldsymbol{\pi}^+ \in \mathcal{P}^{\text{sub}}$ such that

$$\boldsymbol{\pi}_i^+ \in \operatorname*{arg\,max}_{\boldsymbol{\pi}_i' \in \mathcal{F}_i^{\text{markov}}(\boldsymbol{\pi}_{-i})} q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{\pi}_i'(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s}))$$

for all $\boldsymbol{s} \in \mathcal{S}$, and that means $\boldsymbol{\pi}_i^+(\boldsymbol{s}) \in \Phi_i^{\boldsymbol{s}}(\boldsymbol{\pi})$ for all $\boldsymbol{s} \in \mathcal{S}$. Thus, $\Phi_i(\boldsymbol{\pi})$ is also non-empty. Since Cartesian product preserves non-emptiness, compactness, and convexity, we can conclude that $\Phi(\boldsymbol{\pi}) = \bigtimes_{i \in [n]} \Phi_i(\boldsymbol{\pi})$ is non-empty, compact, and convex.

Similarly, fix $i \in [n]$, $\boldsymbol{s} \in \mathcal{S}$, for any $\boldsymbol{\pi} \in \mathcal{P}^{\text{sub}}$, since $\mathcal{X}_i(\boldsymbol{s}, \cdot)$ is continuous (i.e. both upper and lower hemicontinuous), by the Maximum theorem, $\Phi_i^{\boldsymbol{s}}$ is upper hemicontinuous. $\boldsymbol{\pi} \mapsto \bigtimes_{\boldsymbol{s} \in \mathcal{S}} \Phi_i^{\boldsymbol{s}}(\boldsymbol{\pi})$ is upper hemicontinuous as it is a Cartesian product of upper hemicontinuous correspondence, and consequently, $\boldsymbol{\pi} \mapsto \left(\bigtimes_{\boldsymbol{s} \in \mathcal{S}} \Phi_i^{\boldsymbol{s}}(\boldsymbol{\pi})\right) \bigcap \mathcal{P}^{\text{sub}}$ is also upper hemicontinuous. Therefore, Φ is also upper hemicontinuous.

Since $\Phi(\boldsymbol{\pi})$ is non-empty, compact, and convex for any $\boldsymbol{\pi} \in \mathcal{P}^{\text{sub}}$ and Φ is upper hemicontinuous, by the Kakutani-Glicksberg fixed-point theorem (see, Theorem 2.4.1—(Glicksberg, 1952)), Φ admits a fixed point.

Finally, say $\boldsymbol{\pi}^* \in \mathcal{P}^{\text{sub}}$ is a fixed point of Φ, and we want to show that $\boldsymbol{\pi}^*$ is a generalized Markov perfect equilibrium (MPGNE) of $\mathcal{M}$. Since $\boldsymbol{\pi}^* \in \Phi(\boldsymbol{\pi}^*) = \bigtimes_{i \in [n]} \Phi_i(\boldsymbol{\pi}^*)$, we know that for all $i \in [n]$, $\boldsymbol{\pi}_i^*(\boldsymbol{s}) \in \Phi_i^{\boldsymbol{s}}(\boldsymbol{\pi}^*) = \arg\max_{\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{s}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}))} q_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}, \boldsymbol{a}_i, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}))$. We now show that for any $i \in [n]$, for any $\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}_{-i}^*)$, $v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}) \geq v_i^{(\boldsymbol{\pi}_i, \boldsymbol{\pi}_{-i}^*)}(\boldsymbol{s})$ for all $\boldsymbol{s} \in \mathcal{S}$. Take any policy $\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}_{-i}^*)$. Note that $\boldsymbol{\pi}_i$ may not be Markov, so we denote $\{\boldsymbol{\pi}_i(\boldsymbol{h}_{:t})\}_{t \in \mathbb{N}} = \{\boldsymbol{a}_i^{(t)}\}_{t \in \mathbb{N}}$. Then, for all $\boldsymbol{s}^{(0)} \in \mathcal{S}$,

$$
\begin{aligned}
& v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}^{(0)}) \\
& = q_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}^{(0)}, \boldsymbol{\pi}_i^*(\boldsymbol{s}^{(0)}), \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)})) \\
& = \max_{\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{s}^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)}))} q_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}^{(0)}, \boldsymbol{a}_i, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)})) \\
& = \max_{\boldsymbol{a}_i \in \mathcal{X}(\boldsymbol{s}^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)}))} r_i(\boldsymbol{s}^{(0)}, \boldsymbol{a}_i, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)})) + \mathop{\mathbb{E}}_{\boldsymbol{s}^{(1)} \sim \rho(\cdot \mid \boldsymbol{s}^{(0)}, \boldsymbol{a}_i, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)}))} [\gamma v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}^{(1)})] \\
& \geq r_i(\boldsymbol{s}^{(0)}, \boldsymbol{a}_i^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)})) + \mathop{\mathbb{E}}_{\boldsymbol{s}^{(1)} \sim \rho(\cdot \mid \boldsymbol{s}^{(0)}, \boldsymbol{a}_i^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)}))} [\gamma v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}^{(1)})] \qquad (14.5)
\end{aligned}
$$

For any $\boldsymbol{s}^{(0)} \in \mathcal{S}$, define $v_i'(\boldsymbol{s}^{(0)}) \doteq r_i(\boldsymbol{s}^{(0)}, \boldsymbol{a}_i^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)})) + \mathbb{E}_{\boldsymbol{s}^{(1)} \sim \rho(\cdot \mid \boldsymbol{s}^{(0)}, \boldsymbol{a}_i^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)}))}[\gamma v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}^{(1)})]$. Since $v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}) \geq v_i'(\boldsymbol{s})$ for all $i \in [n]$, $\boldsymbol{s} \in \mathcal{S}$, we have for any $\boldsymbol{s}^{(0)} \in \mathcal{S}$

$$
\begin{aligned}
& v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}^{(0)}) \\
& \geq r_i(\boldsymbol{s}^{(0)}, \boldsymbol{a}_i^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)})) + \mathop{\mathbb{E}}_{\boldsymbol{s}^{(1)} \sim \rho(\cdot | \boldsymbol{s}^{(0)}, \boldsymbol{a}_i^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)}))} [\gamma v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}^{(1)})] \\
& \geq r_i(\boldsymbol{s}^{(0)}, \boldsymbol{a}_i^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)})) + \mathop{\mathbb{E}}_{\boldsymbol{s}^{(1)} \sim \rho(\cdot | \boldsymbol{s}^{(0)}, \boldsymbol{a}_i^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)}))} [\gamma v_i'(\boldsymbol{s}^{(1)})] \\
& \geq r_i(\boldsymbol{s}^{(0)}, \boldsymbol{a}_i^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)})) \\
& + \mathop{\mathbb{E}}_{\boldsymbol{s}^{(1)} \sim \rho(\cdot | \boldsymbol{s}^{(0)}, \boldsymbol{a}_i^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)}))} \Bigg[\gamma \Bigg(r_i(\boldsymbol{s}^{(1)}, \boldsymbol{a}_i^{(1)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(1)})) \\
& + \mathop{\mathbb{E}}_{\boldsymbol{s}^{(2)} \sim \rho(\cdot | \boldsymbol{s}^{(1)}, \boldsymbol{a}_i^{(1)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(1)})} [\gamma v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}^{(2)})] \Bigg) \Bigg] \\
& \geq r_i(\boldsymbol{s}^{(0)}, \boldsymbol{a}_i^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)})) \\
& + \mathop{\mathbb{E}}_{\boldsymbol{s}^{(1)} \sim \rho(\cdot | \boldsymbol{s}^{(0)}, \boldsymbol{a}_i^{(0)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(0)}))} \Bigg[\gamma \Bigg(r_i(\boldsymbol{s}^{(1)}, \boldsymbol{a}_i^{(1)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(1)})) \\
& + \mathop{\mathbb{E}}_{\boldsymbol{s}^{(2)} \sim \rho(\cdot | \boldsymbol{s}^{(1)}, \boldsymbol{a}_i^{(1)}, \boldsymbol{\pi}_{-i}^*(\boldsymbol{s}^{(1)})} [\gamma v_i'(\boldsymbol{s}^{(2)})] \Bigg) \Bigg] \\
& \vdots \\
& \geq v_i^{(\boldsymbol{\pi}_i, \boldsymbol{\pi}_{-i}^*)}(\boldsymbol{s})
\end{aligned} \tag{14.6}
$$

where in Equation (14.6), we recursively expand v_i' and eliminate $\boldsymbol{v}^{\boldsymbol{\pi}^*}$ using Equation (14.5). We therefore conclude that for all states $\boldsymbol{s} \in \mathcal{S}$, and for all $i \in [n]$,

$$v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}) \geq \max_{\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}_{-i}^*)} v_i^{(\boldsymbol{\pi}_i, \boldsymbol{\pi}_{-i}^*)}(\boldsymbol{s}).$$

Lemma 12.3.1.

Given a Markov pseudo-game $\mathcal{M}$ for which Assumption 12.2.1 holds, a Markov policy profile $\boldsymbol{\pi}^* \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)$ is a GMPE if and only if $\phi(\boldsymbol{s}, \boldsymbol{\pi}^*) = 0$, for all states $\boldsymbol{s} \in \mathcal{S}$. Similarly, a policy profile $\boldsymbol{\pi}^* \in \mathcal{F}(\boldsymbol{\pi}^*)$ is an GNE if and only if $\varphi(\boldsymbol{\pi}^*) = 0$.

Proof of Lemma 12.3.1

We first prove the result for state exploitability.

($\boldsymbol{\pi}^*$ is a

MPGNE $\implies \phi(\boldsymbol{s}, \boldsymbol{\pi}^*) = 0$ for all $\boldsymbol{s} \in \mathcal{S}$): Suppose that $\boldsymbol{\pi}^*$ is a

MPGNE, i.e., for all players $i \in [n]$, $v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}) \geq \max_{\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}^*_{-i})} v_i^{(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i})}(\boldsymbol{s})$ for all state $\boldsymbol{s} \in \mathcal{S}$. Then, for all state $\boldsymbol{s} \in \mathcal{S}$, we have

$$\forall i \in [n], \quad \max_{\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}^*_{-i})} v_i^{(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i})}(\boldsymbol{s}) - v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}) = 0 \tag{14.7}$$

Summing up across all players $i \in [n]$, we get

$$\phi(\boldsymbol{s}, \boldsymbol{\pi}^*) = \sum_{i \in [n]} \max_{\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}^*_{-i})} v_i^{(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i})}(\boldsymbol{s}) - v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}) = 0 \tag{14.8}$$

($\phi(\boldsymbol{s}, \boldsymbol{\pi}^*) = 0$ for all $\boldsymbol{s} \in \mathcal{S} \implies \boldsymbol{\pi}^*$ is a

MPGNE): Suppose we have $\boldsymbol{\pi}^* \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)$ and $\phi(\boldsymbol{s}, \boldsymbol{\pi}^*) = 0$ for all $\boldsymbol{s} \in \mathcal{S}$. That is, for any $\boldsymbol{s} \in \mathcal{S}$

$$\phi(\boldsymbol{s}, \boldsymbol{\pi}^*) = \sum_{i \in [n]} \max_{\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}^*_{-i})} v_i^{(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i})}(\boldsymbol{s}) - v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}) = 0. \tag{14.9}$$

Since for any $i \in [n]$, $\boldsymbol{\pi}_i^* \in \mathcal{F}_i^{\text{markov}}(\boldsymbol{\pi}^*_{-i})$, $\max_{\boldsymbol{\pi}_i \in \mathcal{F}_i^{\text{markov}}(\boldsymbol{\pi}^*_{-i})} v_i^{(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i})}(\boldsymbol{s}) - v_i^{\boldsymbol{\pi}^*} \geq v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}) - v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}) = 0$. As a result, we must have for all player $i \in [n]$,

$$v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}) = \max_{\boldsymbol{\pi}_i \in \mathcal{F}(\boldsymbol{\pi}^*_{-i})} v_i^{(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i})}(\boldsymbol{s}), \quad \forall \boldsymbol{s} \in \mathcal{S} \tag{14.10}$$

Thus, we can conclude that $\boldsymbol{\pi}^*$ is a

MPGNE.

Then, we can prove results for exploitability in an analogous way.

($\boldsymbol{\pi}^*$ is a GNE $\implies \varphi(\boldsymbol{\pi}^*) = 0$): Suppose that $\boldsymbol{\pi}^*$ is a GNE, i.e., for all players $i \in [n]$, $u_i(\boldsymbol{\pi}^*) \geq \max_{\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}^*_{-i})} u_i(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i})$. Then, we have

$$\forall i \in [n], \quad \max_{\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}^*_{-i})} u_i(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i}) - u_i(\boldsymbol{\pi}^*) = 0 \tag{14.11}$$

Summing up across all players $i \in [n]$, we get

$$\varphi(\boldsymbol{\pi}^*) = \sum_{i \in [n]} \max_{\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}^*_{-i})} u_i(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i}) - u_i(\boldsymbol{\pi}^*) = 0 \tag{14.12}$$

($\varphi(\boldsymbol{s}, \boldsymbol{\pi}^*) = 0 \implies \boldsymbol{\pi}^*$ is a GNE): Suppose we have $\boldsymbol{\pi}^* \in \mathcal{F}(\boldsymbol{\pi}^*)$ and $\varphi(\boldsymbol{\pi}^*) = 0$. That is,

$$\varphi(\boldsymbol{\pi}^*) = \sum_{i \in [n]} \max_{\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}^*_{-i})} u_i(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i}) - u_i(\boldsymbol{\pi}^*) = 0. \tag{14.13}$$

Since for any $i \in [n]$, $\boldsymbol{\pi}^*_i \in \mathcal{F}_i(\boldsymbol{\pi}^*_{-i})$, $\max_{\boldsymbol{\pi}_i \in \mathcal{F}_i(\boldsymbol{\pi}^*_{-i})} u_i(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i}) - u_i(\boldsymbol{\pi}^*) \geq u_i(\boldsymbol{\pi}^*) - u_i(\boldsymbol{\pi}^*) = 0$. As a result, we must have for all player $i \in [n]$,

$$u_i(\boldsymbol{\pi}^*) = \max_{\boldsymbol{\pi}_i \in \mathcal{F}(\boldsymbol{\pi}^*_{-i})} u_i(\boldsymbol{\pi}_i, \boldsymbol{\pi}^*_{-i}) \tag{14.14}$$

Thus, we can conclude that $\boldsymbol{\pi}^*$ is a GNE.

Observation 12.3.1.

Given a Markov pseudo-game $\mathcal{M}$,

$$\min_{\boldsymbol{\pi} \in \mathcal{F}(\boldsymbol{\pi})} \varphi(\boldsymbol{\pi}) = \min_{\boldsymbol{\pi} \in \mathcal{F}(\boldsymbol{\pi})} \max_{\boldsymbol{\pi}' \in \mathcal{F}^{\mathrm{markov}}(\boldsymbol{\pi})} \Psi(\boldsymbol{\pi}, \boldsymbol{\pi}') \ . \tag{12.3}$$

Proof

The per-player maximum operator can be pulled out of the sum in the definition of state-exploitability, because the ith player's best-response policy is independent of the other players' best-response policies, given a fixed policy profile $\boldsymbol{\pi}$:

$$\forall\ \boldsymbol{s} \in \mathcal{S}, \quad \phi(\boldsymbol{s}, \boldsymbol{\pi}) = \sum_{i \in [n]} \max_{\boldsymbol{\pi}'_i \in \mathcal{F}_i^{\mathrm{markov}}(\boldsymbol{\pi}_{-i})} v_i^{(\boldsymbol{\pi}'_i, \boldsymbol{\pi}_{-i})}(\boldsymbol{s}) - v_i^{\boldsymbol{\pi}}(\boldsymbol{s}) \tag{14.15}$$

$$= \max_{\boldsymbol{\pi}' \in \mathcal{F}^{\mathrm{markov}}(\boldsymbol{\pi})} \sum_{i \in [n]} v_i^{(\boldsymbol{\pi}'_i, \boldsymbol{\pi}_{-i})}(\boldsymbol{s}) - v_i^{\boldsymbol{\pi}}(\boldsymbol{s}) \tag{14.16}$$

$$= \max_{\boldsymbol{\pi}' \in \mathcal{F}^{\mathrm{markov}}(\boldsymbol{\pi})} \psi(\boldsymbol{s}, \boldsymbol{\pi}, \boldsymbol{\pi}') \tag{14.17}$$

The argument is analogous for exploitability:

$$\varphi(\boldsymbol{\pi}) = \sum_{i \in [n]} \max_{\boldsymbol{\pi}'_i \in \mathcal{F}_i^{\mathrm{markov}}(\boldsymbol{\pi}_{-i})} u_i(\boldsymbol{\pi}'_i, \boldsymbol{\pi}_{-i}) - u_i(\boldsymbol{\pi}) \tag{14.18}$$

$$= \max_{\boldsymbol{\pi}' \in \mathcal{F}^{\mathrm{markov}}(\boldsymbol{\pi})} \sum_{i \in [n]} u_i(\boldsymbol{\pi}'_i, \boldsymbol{\pi}_{-i}) - u_i(\boldsymbol{\pi}) \tag{14.19}$$

$$= \max_{\boldsymbol{\pi}' \in \mathcal{F}(\boldsymbol{\pi})} \Psi(\boldsymbol{\pi}, \boldsymbol{\pi}') \tag{14.20}$$

Lemma 12.3.3.

Given a Markov pseudo-game $\mathcal{M}$, for $\boldsymbol{\omega} \in \mathbb{R}^{\Omega}$, suppose that $\phi(\boldsymbol{s}, \cdot)$ is differentiable at $\boldsymbol{\omega}$ for all $\boldsymbol{s} \in \mathcal{S}$. If $\|\nabla_{\boldsymbol{\omega}} \varphi(\boldsymbol{\omega})\| = 0$, then, for all states $\boldsymbol{s} \in \mathcal{S}$, $\|\nabla_{\boldsymbol{\omega}} \phi(\boldsymbol{s}, \boldsymbol{\omega})\| = 0$ μ-almost surely, i.e., $\mu(\{\boldsymbol{s} \in \mathcal{S} \mid \|\nabla_{\boldsymbol{\omega}} \phi(\boldsymbol{s}, \boldsymbol{\omega})\| = 0\}) = 1$. Moreover, for any $\varepsilon > 0$ and $\delta \in [0, 1]$, if $\operatorname{supp}(\mu) = \mathcal{S}$ and $\|\nabla_{\boldsymbol{\omega}} \varphi(\boldsymbol{\omega})\| \leq \varepsilon$, then with probability at least $1 - \delta$, $\|\nabla_{\boldsymbol{\omega}} \phi(\boldsymbol{s}, \boldsymbol{\omega})\| \leq \varepsilon/\delta$.

Proof

First, using Jensen's inequality, by the convexity of the 2-norm $\|\cdot\|$, we have:

$$\begin{aligned}
\mathop{\mathbb{E}}_{\boldsymbol{s} \sim \mu} \left[\|\nabla_{\boldsymbol{\omega}} \phi(\boldsymbol{s}, \boldsymbol{\omega})\| \right] &\leq \left\| \mathop{\mathbb{E}}_{\boldsymbol{s} \sim \mu} \left[\nabla_{\boldsymbol{\omega}} \phi(\boldsymbol{s}, \boldsymbol{\omega}) \right] \right\| \\
&= \left\| \nabla_{\boldsymbol{\omega}} \mathop{\mathbb{E}}_{\boldsymbol{s} \sim \mu} [\phi(\boldsymbol{s}, \boldsymbol{\omega})] \right\| \\
&= \|\nabla_{\boldsymbol{\omega}} \varphi(\boldsymbol{\omega})\| \quad .
\end{aligned}$$

The first claim follows directly from the fact that for all $\boldsymbol{s} \in \mathcal{S}$, $\|\nabla_{\boldsymbol{\omega}} \varphi(\boldsymbol{s}, \boldsymbol{\omega})\| \geq 0$, and hence for the expectation $\mathbb{E}_{\boldsymbol{s} \sim \mu} \left[\|\nabla_{\boldsymbol{\omega}} \varphi(\boldsymbol{s}, \boldsymbol{\omega})\| \right]$ to be equal to 0, its value should be equal to zero on a set of measure 1.

Now, for the second part, by Markov's inequality, we have: $\mathbb{P}\left(\|\nabla_{\boldsymbol{\omega}} \phi(\boldsymbol{s}, \boldsymbol{\omega})\| \geq \varepsilon/\delta \right) \leq \frac{\mathbb{E}_{\boldsymbol{s} \sim \mu} \left[\|\nabla_{\boldsymbol{\omega}} \phi(\boldsymbol{s}, \boldsymbol{\pi})\| \right]}{\varepsilon/\delta} \leq \frac{\varepsilon}{\varepsilon/\delta} = \delta$.

Lemma 12.3.4.

Let $\mathcal{M}$ be a Markov pseudo-game with initial state distribution μ. Given policy parameter $\boldsymbol{\omega} \in \mathbb{R}^{\Omega}$ and arbitrary state distribution $\upsilon \in \Delta(\mathcal{S})$, suppose that both $\phi(\mu, \cdot)$ and $\phi(\upsilon, \cdot)$ are differentiable at $\boldsymbol{\omega}$, then we have: $\|\nabla \phi(\upsilon, \boldsymbol{\omega})\| \leq C_{br}(\boldsymbol{\pi}(\cdot; \boldsymbol{\omega}^*), \mu, \upsilon) \|\nabla \varphi(\boldsymbol{\omega})\|$.

Proof

In this proof, for any $i \in [n]$, we define $\boldsymbol{\sigma}_i(\boldsymbol{\omega}) = \boldsymbol{\rho}_i(\cdot, \boldsymbol{\pi}(\cdot; \boldsymbol{\omega}); \boldsymbol{\sigma})$ as player i's policy in the policy profile $\boldsymbol{\sigma}(\boldsymbol{\omega}) = \boldsymbol{\rho}(\cdot, \boldsymbol{\pi}(\cdot; \boldsymbol{\omega}); \boldsymbol{\sigma})$. Similarly, we define $\boldsymbol{\omega}_i = \boldsymbol{\pi}_i(\cdot; \boldsymbol{\omega})$ as player i's policy in the policy profile $\boldsymbol{\omega} = \boldsymbol{\pi}(\cdot; \boldsymbol{\omega})$.

Given a policy parametrization scheme $(\boldsymbol{\pi}, \boldsymbol{\rho}, \mathbb{R}^\Omega, \mathbb{R}^\Sigma)$, consider any two parameters $\boldsymbol{\omega} \in \mathbb{R}^\Omega, \boldsymbol{\sigma} \in \mathbb{R}^\Sigma$, and any two initial state distributions $\mu, \upsilon \in \Delta(\mathcal{S})$, we know that

$$\left\|\nabla_{\boldsymbol{\omega}} \psi(\upsilon, \boldsymbol{\omega}, \boldsymbol{\sigma})\right\| \tag{14.21}$$

$$= \left\|\nabla_{\boldsymbol{\omega}} \sum_{i\in[n]} u_i(\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i}) - u_i(\boldsymbol{\omega})\right\| \tag{14.22}$$

$$= \left\|\sum_{i\in[n]} \nabla_{\boldsymbol{\omega}} (u_i(\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i}) - u_i(\boldsymbol{\omega}))\right\| \tag{14.23}$$

$$= \left\|\sum_{i\in[n]} \nabla_{\boldsymbol{\omega}} \left[\mathbb{E}_{\substack{\boldsymbol{s}'\sim\delta_\upsilon^{(\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i})} \\ \boldsymbol{s}\sim\delta_\upsilon^{\boldsymbol{\omega}}}} \left[r_i(\boldsymbol{s}', \boldsymbol{\rho}_i(\boldsymbol{s}', \boldsymbol{\pi}(\boldsymbol{s};\boldsymbol{\omega});\boldsymbol{\sigma}), \boldsymbol{\pi}_{-i}(\boldsymbol{s}';\boldsymbol{\omega}) - r_i(\boldsymbol{s}, \boldsymbol{\pi}(\boldsymbol{s};\boldsymbol{\omega}))\right]\right]\right\| \tag{14.24}$$

$$= \Big\| \sum_{i\in[n]} \mathbb{E}_{\substack{\boldsymbol{s}'\sim\delta_\upsilon^{(\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i})} \\ \boldsymbol{s}\sim\delta_\upsilon^{\boldsymbol{\omega}}}} \Big[\nabla_{\boldsymbol{a}_{-i}} q_i^{\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i}}(\boldsymbol{s}', \boldsymbol{\rho}_i(\boldsymbol{s}', \boldsymbol{\pi}(\boldsymbol{s}';\boldsymbol{\omega});\boldsymbol{\sigma}), \boldsymbol{\pi}_{-i}(\boldsymbol{s}';\boldsymbol{\omega}))\nabla_{\boldsymbol{\omega}} \left(\boldsymbol{\rho}_i(\boldsymbol{s}', \boldsymbol{\pi}_{-i}(\boldsymbol{s}';\boldsymbol{\omega});\boldsymbol{\omega}), \boldsymbol{\pi}(\boldsymbol{s}';\boldsymbol{\omega})\right)$$

$$- \nabla_{\boldsymbol{a}} q_i^{\boldsymbol{\omega}}(\boldsymbol{s}, \boldsymbol{\pi}(\boldsymbol{s};\boldsymbol{\omega}))\nabla_{\boldsymbol{\omega}} \boldsymbol{\pi}(\boldsymbol{s};\boldsymbol{\omega})\Big]\Big\| \tag{14.25}$$

$$\leq \max_{i\in[n]} \max_{\boldsymbol{s}',\boldsymbol{s}\in\mathcal{S}} \frac{\delta_\upsilon^{(\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i})}(\boldsymbol{s}')\delta_\upsilon^{\boldsymbol{\omega}}(\boldsymbol{s})}{\delta_\mu^{(\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i})}(\boldsymbol{s}')\delta_\mu^{\boldsymbol{\omega}}(\boldsymbol{s})} \Big\| \mathbb{E}_{\substack{\boldsymbol{s}'\sim\delta_\mu^{(\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i})} \\ \boldsymbol{s}\sim\delta_\mu^{\boldsymbol{\omega}}}} \Big[\nabla_{\boldsymbol{a}_{-i}} q_i^{\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i}}(\boldsymbol{s}', \boldsymbol{\rho}_i(\boldsymbol{s}', \boldsymbol{\pi}(\boldsymbol{s}';\boldsymbol{\omega});\boldsymbol{\sigma}), \boldsymbol{\pi}_{-i}(\boldsymbol{s}';\boldsymbol{\omega}))$$

$$\nabla_{\boldsymbol{\omega}} \left(\boldsymbol{\rho}_i(\boldsymbol{s}', \boldsymbol{\pi}_{-i}(\boldsymbol{s}';\boldsymbol{\omega});\boldsymbol{\omega}), \boldsymbol{\pi}(\boldsymbol{s}';\boldsymbol{\omega})\right) - \nabla_{\boldsymbol{a}} q_i^{\boldsymbol{\omega}}(\boldsymbol{s}, \boldsymbol{\pi}(\boldsymbol{s};\boldsymbol{\omega}))\nabla_{\boldsymbol{\omega}} \boldsymbol{\pi}(\boldsymbol{s};\boldsymbol{\omega})\Big]\Big\| \tag{14.26}$$

$$\leq \max_{i\in[n]} \max_{\boldsymbol{s}',\boldsymbol{s}\in\mathcal{S}} \frac{\delta_\upsilon^{(\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i})}(\boldsymbol{s}')\delta_\upsilon^{\boldsymbol{\omega}}(\boldsymbol{s})}{\delta_\mu^{(\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i})}(\boldsymbol{s}')\delta_\mu^{\boldsymbol{\omega}}(\boldsymbol{s})} \left\|\nabla_{\boldsymbol{\omega}} \left[v_i^{\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i}}(\mu) - v_i^{\boldsymbol{\omega}}(\mu)\right]\right\| \tag{14.27}$$

$$\leq \left(\frac{1}{1-\gamma}\right)^2 \max_{i\in[n]} \max_{\boldsymbol{s}',\boldsymbol{s}\in\mathcal{S}} \frac{\delta_\upsilon^{(\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i})}(\boldsymbol{s}')\delta_\upsilon^{\boldsymbol{\omega}}(\boldsymbol{s})}{\mu(\boldsymbol{s}')\mu(\boldsymbol{s})} \left\|\nabla_{\boldsymbol{\omega}} \psi(\mu, \boldsymbol{\omega}, \boldsymbol{\sigma})\right\| \tag{14.28}$$

$$= \left(\frac{1}{1-\gamma}\right)^2 \max_{i\in[n]} \left\|\frac{\delta_\upsilon^{(\boldsymbol{\sigma}_i(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i})}}{\mu}\right\|_\infty \left\|\frac{\delta_\mu^{\boldsymbol{\omega}}}{\mu}\right\|_\infty \left\|\nabla_{\boldsymbol{\omega}} \psi(\mu, \boldsymbol{\omega}, \boldsymbol{\sigma})\right\| \tag{14.29}$$

where Equation (14.25) and Equation (14.27) are obtained by deterministic policy gradient theorem (Silver et al., 2014), and Equation (14.28) is due to the fact that $\delta_\mu^{\boldsymbol{\omega}}(\boldsymbol{s}) \geq (1-\gamma)\mu(\boldsymbol{s})$ for any $\boldsymbol{\pi} \in \mathcal{P}$, $\boldsymbol{s} \in \mathcal{S}$.

Given condition (1) of Assumption 12.3.3, fix any $\boldsymbol{\omega} \in \mathbb{R}^\Omega$, there exists $\boldsymbol{\sigma}^* \in \mathbb{R}^\Sigma$ s.t. for all $i \in [n]$, $\boldsymbol{s} \in \mathcal{S}$:

$$q_i^{\boldsymbol{\omega}}(\boldsymbol{s}, \boldsymbol{\rho}_i(\boldsymbol{s}, \boldsymbol{\pi}(\boldsymbol{s};\boldsymbol{\omega});\boldsymbol{\sigma}^*), \boldsymbol{\pi}_{-i}(\boldsymbol{s};\boldsymbol{\omega})) = \max_{\boldsymbol{\pi}_i' \in \mathcal{F}_i(\boldsymbol{\pi}(\cdot;\boldsymbol{\omega}))} q_i^{\boldsymbol{\omega}}(\boldsymbol{s}, \boldsymbol{\pi}_i'(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s};\boldsymbol{\omega})) \ .$$

Thus, $\phi(\boldsymbol{s},\boldsymbol{\omega}) = \psi(\boldsymbol{s},\boldsymbol{\omega},\boldsymbol{\sigma}^*)$ for all $\boldsymbol{s} \in \mathcal{S}$. Hence, plugging in the optimal best-response policy $\boldsymbol{\sigma} = \boldsymbol{\sigma}^*$, we obtain that

$$\|\nabla_{\boldsymbol{\omega}}\phi(\upsilon,\boldsymbol{\omega})\| \leq \left(\frac{1}{1-\gamma}\right)^2 \max_{i\in[n]} \left\|\frac{\delta_\upsilon^{(\boldsymbol{\sigma}_i^*(\boldsymbol{\omega}),\boldsymbol{\omega}_{-i})}}{\mu}\right\|_\infty \left\|\frac{\delta_\mu^{\boldsymbol{\omega}}}{\mu}\right\|_\infty \|\nabla_{\boldsymbol{\omega}}\phi(\mu,\boldsymbol{\omega})\| \tag{14.30}$$

$$\leq \left(\frac{1}{1-\gamma}\right)^2 \max_{i\in[n]} \max_{\boldsymbol{\pi}_i'\in\Phi_i(\boldsymbol{\pi}_{-i}(\cdot;\boldsymbol{\omega}))} \left\|\frac{\delta_\upsilon^{(\boldsymbol{\pi}_i',\boldsymbol{\pi}_{-i}(\cdot;\boldsymbol{\omega}))}}{\mu}\right\|_\infty \left\|\frac{\delta_\mu^{\boldsymbol{\omega}}}{\mu}\right\|_\infty \|\nabla_{\boldsymbol{\omega}}\phi(\mu,\boldsymbol{\omega})\| \tag{14.31}$$

where eq. (14.31) is due to the fact that $\boldsymbol{\sigma}_i^*(\boldsymbol{\omega}) \in \Phi_i(\boldsymbol{\pi}_{-i}(\cdot;\boldsymbol{\omega}))$.

Lemma 12.3.2.

Given a Markov pseudo-game $\mathcal{M}$,

$$\min_{\boldsymbol{\pi}\in\mathcal{F}(\boldsymbol{\pi})} \max_{\boldsymbol{\pi}'\in\mathcal{F}^{\text{markov}}(\boldsymbol{\pi})} \Psi(\boldsymbol{\pi},\boldsymbol{\pi}') = \min_{\boldsymbol{\pi}\in\mathcal{F}(\boldsymbol{\pi})} \max_{\boldsymbol{\rho}\in\mathcal{R}} \Psi(\boldsymbol{\pi},\boldsymbol{\rho}(\cdot,\boldsymbol{\pi}(\cdot))) \ . \tag{12.4}$$

Proof

Fix $\boldsymbol{\pi}^* \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)$. We want to show that

$$\max_{\boldsymbol{\pi}'\in\mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)} \varphi(\boldsymbol{\pi}^*,\boldsymbol{\pi}') = \max_{\boldsymbol{\rho}\in\mathcal{R}} \varphi(\boldsymbol{\pi}^*,\boldsymbol{\rho}(\cdot,\boldsymbol{\pi}(\cdot))) \ .$$

Define $\mathcal{P}^{\mathcal{R},\boldsymbol{\pi}^*} \doteq \{\boldsymbol{\pi} : \boldsymbol{s} \mapsto \boldsymbol{\rho}(\boldsymbol{s},\boldsymbol{\pi}^*(\boldsymbol{s})) \mid \boldsymbol{\rho} \in \mathcal{R}\} \subseteq \mathcal{P}^{\text{markov}}$.

First, for all $\boldsymbol{\pi}' \in \mathcal{P}^{\mathcal{R},\boldsymbol{\pi}^*}$, $\boldsymbol{\pi}'(\boldsymbol{s}) = \boldsymbol{\rho}(\boldsymbol{s},\boldsymbol{\pi}^*(\boldsymbol{s})) \in \mathcal{X}(\boldsymbol{s},\boldsymbol{\pi}^*(\boldsymbol{s}))$, for all $\boldsymbol{s} \in \mathcal{S}$, by the definition of $\mathcal{R}$. Thus, $\boldsymbol{\pi}' \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*) = \{\boldsymbol{\pi} \in \mathcal{P}^{\text{markov}} \mid \forall \boldsymbol{s} \in \mathcal{S}, \boldsymbol{\pi}(\boldsymbol{s}) \in \mathcal{X}(\boldsymbol{s},\boldsymbol{\pi}^*(\boldsymbol{s}))\}$.

Therefore, $\mathcal{P}^{\mathcal{R},\boldsymbol{\pi}^*} \subseteq \mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)$, which implies that $\max_{\boldsymbol{\pi}'\in\mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)} \varphi(\boldsymbol{\pi}^*,\boldsymbol{\pi}') \geq \max_{\boldsymbol{\pi}'\in\mathcal{P}^{\mathcal{R},\boldsymbol{\pi}^*}} \varphi(\boldsymbol{\pi}^*,\boldsymbol{\pi}') = \max_{\boldsymbol{\rho}\in\mathcal{R}} \varphi(\boldsymbol{\pi}^*,\boldsymbol{\rho}(\cdot,\boldsymbol{\pi}(\cdot)))$.

Moreover, for all $\boldsymbol{\pi}' \in \mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)$, $\boldsymbol{\pi}'(\boldsymbol{s}) \in \mathcal{X}(\boldsymbol{s},\boldsymbol{\pi}^*(\boldsymbol{s}))$, for all $\boldsymbol{s} \in \mathcal{S}$, by the definition of $\mathcal{F}^{\text{markov}}$. Define $\boldsymbol{\rho}'$ such that for all $\boldsymbol{s} \in \mathcal{S}$, $\boldsymbol{\rho}'(\boldsymbol{s},\boldsymbol{a}) = \boldsymbol{\pi}'(\boldsymbol{s})$ if $\boldsymbol{a} = \boldsymbol{\pi}^*(\boldsymbol{s})$, and $\boldsymbol{\rho}'(\boldsymbol{s},\boldsymbol{a}) = \boldsymbol{a}'$ for some $\boldsymbol{a}' \in \mathcal{X}(\boldsymbol{s},\boldsymbol{a})$ otherwise. Note that $\boldsymbol{\rho}' \in \mathcal{R}$, since $\forall(\boldsymbol{s},\boldsymbol{a}) \in \mathcal{S}\times\mathcal{A},\ \boldsymbol{\rho}(\boldsymbol{s},\boldsymbol{a}) \in \mathcal{X}(\boldsymbol{s},\boldsymbol{a})$. Thus, as $\boldsymbol{\pi}'(\boldsymbol{s}) = \boldsymbol{\rho}'(\boldsymbol{s},\boldsymbol{\pi}^*(\boldsymbol{s}))$, for all $\boldsymbol{s} \in \mathcal{S}$, it follows that $\boldsymbol{\pi}' \in \mathcal{P}^{\mathcal{R},\boldsymbol{\pi}^*}$. Therefore, $\mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*) \subseteq \mathcal{P}^{\mathcal{R},\boldsymbol{\pi}^*}$, which implies that $\max_{\boldsymbol{\pi}'\in\mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)} \varphi(\boldsymbol{\pi}^*,\boldsymbol{\pi}') \leq \max_{\boldsymbol{\pi}'\in\mathcal{P}^{\mathcal{R},\boldsymbol{\pi}^*}} \varphi(\boldsymbol{\pi}^*,\boldsymbol{\pi}') = \max_{\boldsymbol{\rho}\in\mathcal{R}} \varphi(\boldsymbol{\pi}^*,\boldsymbol{\rho}(\cdot,\boldsymbol{\pi}(\cdot)))$.

Finally, we conclude that $\max_{\boldsymbol{\pi}'\in\mathcal{F}^{\text{markov}}(\boldsymbol{\pi}^*)} \varphi(\boldsymbol{\pi}^*,\boldsymbol{\pi}') = \max_{\boldsymbol{\rho}\in\mathcal{R}} \varphi(\boldsymbol{\pi}^*,\boldsymbol{\rho}(\cdot,\boldsymbol{\pi}(\cdot)))$.

Theorem 12.3.1.

Given a Markov pseudo-game $\mathcal{M}$, and a parameterization scheme $(\boldsymbol{\pi}, \boldsymbol{\rho}, \mathbb{R}^\Omega, \mathbb{R}^\Sigma)$, suppose Assumption 12.2.1, 12.3.2, and 12.3.3 hold. For any $\delta > 0$, set $\varepsilon = \delta \| C_{br}(\cdot, \mu, \cdot) \|_\infty^{-1}$. If Algorithm 11 is run with inputs that satisfy, $\eta_{\boldsymbol{\omega}}, \eta_{\boldsymbol{\sigma}} \asymp \text{poly}(\varepsilon, \|\partial \delta_\mu^{\boldsymbol{\pi}^*} / \partial \mu\|_\infty, \frac{1}{1-\gamma}, \ell_{\nabla \Psi}^{-1}, \ell_\Psi^{-1})$, then for some $T \in \text{poly}\left(\varepsilon^{-1}, (1-\gamma)^{-1}, \|\partial \delta_\mu^{\boldsymbol{\pi}^*} / \partial \mu\|_\infty, \ell_{\nabla \Psi}, \ell_\Psi, \text{diam}(\mathbb{R}^\Omega \times \mathbb{R}^\Sigma), \eta_{\boldsymbol{\omega}}^{-1}\right)$, there exists $\boldsymbol{\omega}_{\text{best}}^{(T)} = \boldsymbol{\omega}^{(k)}$ with $k \leq T$ that is a $(\varepsilon, \varepsilon/2\ell_\Psi)$-stationary point of the exploitability, i.e., there exists $\boldsymbol{\omega}^* \in \mathbb{R}^\Omega$ s.t. $\|\boldsymbol{\omega}_{\text{best}}^{(T)} - \boldsymbol{\omega}^*\| \leq \varepsilon/2\ell_\Psi$ and $\min_{\boldsymbol{h} \in \mathcal{D}\varphi(\boldsymbol{\omega}^*)} \|\boldsymbol{h}\| \leq \varepsilon$.

Further, for any arbitrary state distribution $\upsilon \in \Delta(\mathcal{S})$, if $\phi(\upsilon, \cdot)$ is differentiable at $\boldsymbol{\omega}^*$, $\|\nabla_{\boldsymbol{\omega}} \varphi(\upsilon, \boldsymbol{\omega}^*)\| \leq \delta$, i.e., $\boldsymbol{\omega}_{\text{best}}^{(T)}$ is a (ε, δ)-stationary point for the expected exploitability $\phi(\upsilon, \cdot)$.

Proof

As is common in the optimization literature (see, for instance, Davis et al. (2018)), we consider the Moreau envelope of the exploitability, which we simply call the **Moreau exploitability**, i.e.,

$$\tilde{\varphi}(\boldsymbol{\omega}) \doteq \min_{\boldsymbol{\omega}' \in \mathbb{R}^\Omega} \left\{ \varphi(\boldsymbol{\omega}') + \ell_{\nabla \psi} \left\| \boldsymbol{\omega} - \boldsymbol{\omega}' \right\|^2 \right\} \quad .$$

Similarly, we also consider the **state Moreau exploitability**, i.e., the Moreau envelope of the state exploitability:

$$\tilde{\phi}(\boldsymbol{s}, \boldsymbol{\omega}) \doteq \min_{\boldsymbol{\omega}' \in \mathbb{R}^\Omega} \left\{ \phi(\boldsymbol{s}, \boldsymbol{\omega}') + \ell_{\nabla \psi} \left\| \boldsymbol{\omega} - \boldsymbol{\omega}' \right\|^2 \right\} \quad .$$

We recall that in these definitions, by our notational convention, $\ell_{\nabla \psi} \geq 0$, refers to the Lipschitz-smoothness constants of the state exploitability which in this case we take to be the largest across all states, i.e., for all $\boldsymbol{s} \in \mathcal{S}$, $(\boldsymbol{\omega}, \boldsymbol{\sigma}) \mapsto \psi(\boldsymbol{s}, \boldsymbol{\omega}, \boldsymbol{\sigma})$ is $\ell_{\nabla \psi}$-Lipschitz-smooth, respectively, and which we note is guaranteed to exist under Assumption 12.3.2. Further, we note that since $\Psi(\boldsymbol{\omega}, \boldsymbol{\sigma}) = \mathbb{E}_{\boldsymbol{s} \sim \mu} [\psi(\boldsymbol{s}, \boldsymbol{\omega}, \boldsymbol{\sigma})]$ is a weighted average of ψ, $(\boldsymbol{\omega}, \boldsymbol{\sigma}) \mapsto \Psi(\boldsymbol{\omega}, \boldsymbol{\sigma})$ is also $\ell_{\nabla \psi}$-Lipschitz-smooth.

We invoke Theorem 2 of Daskalakis et al. (2020a). Although their result is stated for gradient-dominated-gradient-dominated functions, their proof applies in the more general case of non-convex-gradient-dominated functions.

First, Assumption 12.3.2 guarantees that the cumulative regret Ψ is Lipschitz-smooth w.r.t. $(\boldsymbol{\omega}, \boldsymbol{\sigma})$. Moreover, under Assumption 12.3.2, which guarantees that $\boldsymbol{\sigma} \mapsto q_i^{\boldsymbol{\omega}'}(\boldsymbol{s}, \boldsymbol{\rho}_i(\boldsymbol{s}, \boldsymbol{\pi}_{-i}(\boldsymbol{s}; \boldsymbol{\omega}); \boldsymbol{\sigma}), \boldsymbol{\pi}_{-i}(\boldsymbol{s}; \boldsymbol{\omega}))$ is continuously differentiable for all $\boldsymbol{s} \in \mathcal{S}$ and $\boldsymbol{\omega}, \boldsymbol{\omega}' \in \mathbb{R}^{\Omega}$, and Assumption 12.3.3, we have that Ψ is $\left(\left\| {\partial \delta_{\mu}^{\boldsymbol{\pi}^*}}/{\partial \mu} \right\|_{\infty} / 1-\gamma\right)$-gradient-dominated in $\boldsymbol{\sigma}$, for all $\boldsymbol{\omega} \in \mathbb{R}^{\Omega}$, by Theorems 2 and 4 of Bhandari and Russo (2019). Finally, under Assumption 12.3.2, since the policy, the reward function, and the transition probability function are all Lipschitz-continuous, $\widehat{u}$, $\widehat{\Psi}$, and hence $\widehat{G}$ are also Lipschitz-continuous, since $\mathcal{S}$ and $\mathcal{A}$ are compact. Their variance must therefore be bounded, i.e., there exists $\varsigma_{\boldsymbol{\omega}}, \varsigma_{\boldsymbol{\sigma}} \in \mathbb{R}$ s.t. $\mathbb{E}_{\boldsymbol{h},\boldsymbol{h}'}[\widehat{G_{\boldsymbol{\omega}}}(\boldsymbol{\omega}, \boldsymbol{\sigma}; \boldsymbol{h}, \boldsymbol{h}') - \nabla_{\boldsymbol{\omega}} \Psi(\boldsymbol{\omega}, \boldsymbol{\sigma}; \boldsymbol{h}, \boldsymbol{h}')] \leq \varsigma_{\boldsymbol{\omega}}$ and $\mathbb{E}_{\boldsymbol{h},\boldsymbol{h}'}[\widehat{G_{\boldsymbol{\sigma}}}(\boldsymbol{\omega}, \boldsymbol{\sigma}; \boldsymbol{h}, \boldsymbol{h}') - \nabla_{\boldsymbol{\sigma}} \Psi(\boldsymbol{\omega}, \boldsymbol{\sigma}; \boldsymbol{h}, \boldsymbol{h}')] \leq \varsigma_{\boldsymbol{\sigma}}$.

Hence, under our assumptions, the assumptions of Theorem 2 of Daskalakis et al. are satisfied. Therefore, $1/T+1 \sum_{t=0}^{T} \|\nabla \widetilde{\varphi}(\boldsymbol{\omega}^{(t)})\| \leq \varepsilon$. Taking a minimum across all $t \in [T]$, we conclude $\left\|\nabla \widetilde{\varphi}(\boldsymbol{\omega}_{\text{best}}^{(T)})\right\| \leq \varepsilon$.

Then, by the Lemma 3.7 of (Lin et al., 2020), there exists some $\boldsymbol{\omega}^* \in \mathbb{R}^{\Omega}$ such that $\|\boldsymbol{\omega}_{\text{best}}^{(T)} - \boldsymbol{\omega}^*\| \leq \frac{\varepsilon}{2\ell_{\Psi}}$ and $\boldsymbol{\omega}^* \in \mathbb{R}^{\Omega}_{\varepsilon} \doteq \{\boldsymbol{\omega} \in \mathbb{R}^{\Omega} \mid \exists \alpha \in \mathcal{D}\varphi(\boldsymbol{\omega}), \|\alpha\| \leq \varepsilon\}$. That is, $\boldsymbol{\omega}_{\text{best}}^{(T)}$ is a $(\varepsilon, \frac{\varepsilon}{2\ell_{\Psi}})$-stationary point of φ.

Furthermore, if we assume that $\phi(\delta, \cdot)$ is differentiable at $\boldsymbol{\omega}^*$ for any state distribution $\delta \in \Delta(\mathcal{S})$, φ is also differentiable at $\boldsymbol{\omega}^*$. Hence, by the proof of Lemma 12.3.4, we know that for any state distribution $\upsilon \in \Delta(\mathcal{S})$,

$$\|\nabla_{\boldsymbol{\omega}} \phi(\upsilon, \boldsymbol{\omega})\| \leq \max_{\boldsymbol{\sigma}^* \in \arg\max_{\boldsymbol{\sigma} \in \mathbb{R}^{\Sigma}} \psi(\upsilon, \boldsymbol{\omega}, \boldsymbol{\sigma})} \|\nabla_{\boldsymbol{\omega}} \psi(\upsilon, \boldsymbol{\omega}, \boldsymbol{\sigma}^*)\| \tag{14.32}$$

$$\leq \max_{i \in [n]} \max_{\boldsymbol{\sigma}^* \in \arg\max_{\boldsymbol{\sigma} \in \mathbb{R}^{\Sigma}} \psi(\upsilon, \boldsymbol{\omega}, \boldsymbol{\sigma})} \tag{14.33}$$

$$\left(\frac{1}{1-\gamma}\right)^2 \left\| \frac{\delta_{\upsilon}^{\boldsymbol{\sigma}_i^*(\boldsymbol{\omega}), \boldsymbol{\omega}_{-i}}}{\mu} \right\|_{\infty} \left\| \frac{\delta_{\upsilon}^{\boldsymbol{\omega}}}{\mu} \right\|_{\infty} \|\nabla_{\boldsymbol{\omega}} \Psi(\boldsymbol{\omega}, \boldsymbol{\sigma}^*)\| \tag{14.34}$$

$$= C_{br}(\boldsymbol{\omega}, \mu, \upsilon) \|\nabla_{\boldsymbol{\omega}} \Psi(\boldsymbol{\omega}, \boldsymbol{\sigma}^*)\| \tag{14.35}$$

$$\frac{1}{C_{br}(\boldsymbol{\omega}, \mu, \upsilon)} \|\nabla_{\boldsymbol{\omega}} \phi(\upsilon, \boldsymbol{\omega})\| \leq \|\nabla_{\boldsymbol{\omega}} \Psi(\boldsymbol{\omega}, \boldsymbol{\sigma}^*)\| \tag{14.36}$$

Therefore,

$$\boldsymbol{\omega}^* \in \mathbb{R}^{\Omega}_{\varepsilon} \doteq \{\boldsymbol{\omega} \in \mathbb{R}^{\Omega} \mid \exists \alpha \in \mathcal{D}\varphi(\boldsymbol{\omega}), \|\alpha\| \leq \varepsilon\} \tag{14.37}$$

$$\supseteq \{\boldsymbol{\omega} \in \mathbb{R}^{\Omega} \mid \exists \boldsymbol{\sigma}^* \in \arg\max_{\boldsymbol{\omega} \in \mathbb{R}^{\Omega}} \Psi(\boldsymbol{\omega}, \boldsymbol{\sigma}) s.t. \|\nabla_{\boldsymbol{\omega}} \Psi(\boldsymbol{\omega}, \boldsymbol{\sigma}^*)\| \leq \varepsilon\} \tag{14.38}$$

$$\supseteq \{\boldsymbol{\omega} \in \mathbb{R}^{\Omega} \mid {}^{1}/_{C_{br}(\boldsymbol{\omega}, \mu, \upsilon)} \|\nabla_{\boldsymbol{\omega}} \phi(\upsilon, \boldsymbol{\omega})\| \leq \varepsilon\} \tag{14.39}$$

$$= \{\boldsymbol{\omega} \in \mathbb{R}^{\Omega} \mid \|\nabla_{\boldsymbol{\omega}} \phi(\upsilon, \boldsymbol{\omega})\| \leq \delta\} \tag{14.40}$$

Therefore, we can conclude that there exists $\boldsymbol{\omega}^*$ such that $\|\boldsymbol{\omega}^{(T)}_{\text{best}} - \boldsymbol{\omega}^*\| \leq \frac{\varepsilon}{2\ell_{\Psi}}$ and $\|\nabla_{\boldsymbol{\omega}} \phi(\upsilon, \boldsymbol{\omega})\| \leq \delta$ for any υ. Thus, $\boldsymbol{\omega}^{(T)}_{\text{best}}$ is a (ε, δ)-stationary point of $\phi(\upsilon, \cdot)$ for any $\upsilon \in \Delta(\mathcal{S})$.

14.1.2 Omitted Results and Proofs from Section 13.1.3

Theorem 13.1.1.

Consider a Radner economy $\mathcal{I}$. Under Assumption 13.1.1, the set of recursive Radner equilibria of $\mathcal{I}$ is equal to the set of GMPE of the associated Radner Markov pseudo-game $\mathcal{M}$.

Proof

Let $\boldsymbol{\pi}^* = (\boldsymbol{X}^*, \boldsymbol{Y}^*, \boldsymbol{p}^*, \boldsymbol{q}^*) : \mathcal{S} \to \mathcal{X} \times \mathcal{Y} \times \mathcal{P} \times \mathcal{Q}$ be an GMPE of the Radner Markov pseudo-game $\mathcal{M}$ associated with $\mathcal{I}$. We want to show that it is also an RRE of $\mathcal{I}$.

First, we want to show that $\boldsymbol{\pi}^*$ is Markov perfect for all consumers. We can make some easy observations: the state value for the player $i \in [n]$ in the Radner Markov pseudo-game at state $\boldsymbol{s} \in \mathcal{S}$ induced by the policy $\boldsymbol{\pi}^*$

$$v_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}) = \mathop{\mathbb{E}}_{H \sim \nu^{\boldsymbol{\pi}^*}} \left[\sum_{t=0}^{\infty} \gamma^t r'(S^{(t)}, A^{(t)}) \mid S^{(0)} = \boldsymbol{s}) \right] \tag{14.41}$$

$$= \mathop{\mathbb{E}}_{H \sim \nu^{\boldsymbol{\pi}^*}} \left[\sum_{t=0}^{\infty} \gamma^t u_i(\boldsymbol{x}_i^*(S^{(t)}); \Theta_i^{(t)}) \mid S^{(0)} = \boldsymbol{s}) \right] \tag{14.42}$$

is equal to the consumption state value induced by $(\boldsymbol{X}^*, \boldsymbol{Y}^*, \boldsymbol{p}^*, \boldsymbol{q}^*)$

$$v_i^{(\boldsymbol{X}^*,\boldsymbol{Y}^*,\boldsymbol{p}^*,\boldsymbol{q}^*)}(\boldsymbol{s}) \doteq \mathop{\mathbb{E}}_{H\sim\nu^{(\boldsymbol{X}^*,\boldsymbol{Y}^*,\boldsymbol{p}^*,\boldsymbol{q}^*)}} \left[\sum_{t=0}^{\infty} \gamma^t u_i\left(\boldsymbol{x}_i^*(H_{:t}); \Theta^{(t)}\right) \mid S^{(0)} = \boldsymbol{s}\right] \quad . \tag{14.43}$$

as $\boldsymbol{x}_i^*$ is Markov. Since $\boldsymbol{\pi}^*$ is a GMPE, we know that for any $i \in [n]$:

$$(\boldsymbol{x}_i^*, \boldsymbol{y}_i^*) \in \mathop{\arg\max}_{\substack{(\boldsymbol{x}_i,\boldsymbol{y}_i):\mathcal{S}\to\mathcal{X}_i\times\mathcal{Y}_i:\forall\boldsymbol{s}\in\mathcal{S},\\ (\boldsymbol{x}_i,\boldsymbol{y}_i)(\boldsymbol{s})\in\mathcal{B}_i(\boldsymbol{e}_i,\boldsymbol{p}^*(\boldsymbol{s}),\boldsymbol{q}^*(\boldsymbol{s}))}} \left\{v_i^{(\boldsymbol{x}_i,\boldsymbol{x}^*_{-i},\boldsymbol{y}_i,\boldsymbol{y}^*_{-i},\boldsymbol{p}^*,\boldsymbol{q}^*)}(\boldsymbol{s})\right\}$$

for all $\boldsymbol{s} \in \mathcal{S}$, so $(\boldsymbol{X}^*, \boldsymbol{Y}^*, \boldsymbol{p}^*, \boldsymbol{q}^*)$ is Markov perfect.

Next, we want to show that $(\boldsymbol{X}^*, \boldsymbol{Y}^*, \boldsymbol{p}^*, \boldsymbol{q}^*)$ satisfies the Walras's law. First, we show that for any $i \in [n]$, $\boldsymbol{s} \in \mathcal{S}$, $\boldsymbol{x}_i^*(\boldsymbol{s}) \cdot \boldsymbol{p}^*(\boldsymbol{s}) + \boldsymbol{y}_i^*(\boldsymbol{s}) \cdot \boldsymbol{q}^*(\boldsymbol{s}) - \boldsymbol{e}_i \cdot \boldsymbol{p}^*(\boldsymbol{s}) = 0$. By way of contradiction, assume that there exists some $i \in [n]$, $\boldsymbol{s} \in \mathcal{S}$ such that $\boldsymbol{x}_i^*(\boldsymbol{s}) \cdot \boldsymbol{p}^*(\boldsymbol{s}) + \boldsymbol{y}_i^*(\boldsymbol{s}) \cdot \boldsymbol{q}^*(\boldsymbol{s}) - \boldsymbol{e}_i \cdot \boldsymbol{p}^*(\boldsymbol{s}) \neq 0$. Note that $(\boldsymbol{x}_i^*(\boldsymbol{s}), \boldsymbol{y}_i^*(\boldsymbol{s})) \in \mathcal{B}'(\boldsymbol{s}, \boldsymbol{a}_{-i}) = \mathcal{B}(\boldsymbol{e}_i, \boldsymbol{p}^*(\boldsymbol{s}), \boldsymbol{q}^*(\boldsymbol{s})) = \{(\boldsymbol{x}_i, \boldsymbol{y}_i) \in \mathcal{X}_i \times \mathcal{Y}_i \mid \boldsymbol{x}_i \cdot \boldsymbol{p}^*(\boldsymbol{s}) + \boldsymbol{y}_i \cdot \boldsymbol{q}^*(\boldsymbol{s}) \leq \boldsymbol{e}_i \cdot \boldsymbol{p}^*(\boldsymbol{s})\}$, so we must have $\boldsymbol{x}_i^*(\boldsymbol{s}) \cdot \boldsymbol{p}^*(\boldsymbol{s}) + \boldsymbol{y}_i^*(\boldsymbol{s}) \cdot \boldsymbol{q}^*(\boldsymbol{s}) - \boldsymbol{e}_i \cdot \boldsymbol{p}^*(\boldsymbol{s}) < 0$. By the (no saturation) condition of Assumption 13.1.1, there exists $\boldsymbol{x}_i^+ \in \mathcal{X}_i$ s.t. $u_i(\boldsymbol{x}_i^+; \boldsymbol{\theta}_i) > u_i(\boldsymbol{x}_i^*(\boldsymbol{s}); \boldsymbol{\theta}_i)$. Moreover, since $\boldsymbol{x}_i \mapsto u_i(\boldsymbol{x}_i; \boldsymbol{\theta}_i)$ is concave, for any $0 < t < 1$, $u_i(t\boldsymbol{x}_i^+ + (1-t)\boldsymbol{x}_i^*(\boldsymbol{s}); \boldsymbol{\theta}_i) > u_i(\boldsymbol{x}_i^*(\boldsymbol{s}); \boldsymbol{\theta}_i)$. Since $\boldsymbol{x}_i^*(\boldsymbol{s}) \cdot \boldsymbol{p}^*(\boldsymbol{s}) + \boldsymbol{y}_i^*(\boldsymbol{s}) \cdot \boldsymbol{q}^*(\boldsymbol{s}) - \boldsymbol{e}_i \cdot \boldsymbol{p}^*(\boldsymbol{s}) < 0$, we can pick t small enough such that $\boldsymbol{x}_i' = t\boldsymbol{x}_i^+ + (1-t)\boldsymbol{x}_i^*(\boldsymbol{s})$ satisfies $\boldsymbol{x}_i'(\boldsymbol{s}) \cdot \boldsymbol{p}^*(\boldsymbol{s}) + \boldsymbol{y}_i^*(\boldsymbol{s}) \cdot \boldsymbol{q}^*(\boldsymbol{s}) - \boldsymbol{e}_i \cdot \boldsymbol{p}^*(\boldsymbol{s}) \leq 0$ but $\boldsymbol{x}_i' \in \mathcal{X}_i$ s.t. $u_i(\boldsymbol{x}_i^+; \boldsymbol{\theta}_i) > u_i(\boldsymbol{x}_i^*(\boldsymbol{s}); \boldsymbol{\theta}_i)$. Thus,

$$q_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}, \boldsymbol{x}_i', \boldsymbol{x}_{-i}^*(\boldsymbol{s}), \boldsymbol{Y}^*(\boldsymbol{s}), \boldsymbol{p}^*(\boldsymbol{s}), \boldsymbol{q}^*(\boldsymbol{s})) \tag{14.44}$$

$$= r_i'(\boldsymbol{s}, \boldsymbol{x}_i', \boldsymbol{x}_{-i}^*(\boldsymbol{s}), \boldsymbol{Y}^*(\boldsymbol{s}), \boldsymbol{p}^*(\boldsymbol{s}), \boldsymbol{q}^*(\boldsymbol{s})) + \mathop{\mathbb{E}}_{S'\sim\rho(S'|\boldsymbol{s},\boldsymbol{Y}^*(\boldsymbol{s}))}[\gamma v_i^{\boldsymbol{\pi}^*}(S')] \tag{14.45}$$

$$= u_i(\boldsymbol{x}_i'; \boldsymbol{\theta}_i) + \mathop{\mathbb{E}}_{S'\sim\rho(S'|\boldsymbol{s},\boldsymbol{Y}^*(\boldsymbol{s}))}[\gamma v_i^{\boldsymbol{\pi}^*}(S')] \tag{14.46}$$

$$> u_i(\boldsymbol{x}_i^*(\boldsymbol{s}); \boldsymbol{\theta}_i) + \mathop{\mathbb{E}}_{S'\sim\rho(S'|\boldsymbol{s},\boldsymbol{Y}^*(\boldsymbol{s}))}[\gamma v_i^{\boldsymbol{\pi}^*}(S')] \tag{14.47}$$

$$= q_i^{\boldsymbol{\pi}^*}(\boldsymbol{s}, \boldsymbol{X}^*(\boldsymbol{s}), \boldsymbol{Y}^*(\boldsymbol{s}), \boldsymbol{p}^*(\boldsymbol{s}), \boldsymbol{q}^*(\boldsymbol{s})) \tag{14.48}$$

This contradicts that fact that $\boldsymbol{\pi}^*$ is a GMPE since an optimal policy is supposed to be greedy optimal (i.e., maximize the action-value function of each player over its action space at all states) respect to optimal action value function. Thus, we know that for all $i \in [n]$, $\boldsymbol{s} \in \mathcal{S}$, $\boldsymbol{x}_i^*(\boldsymbol{s}) \cdot \boldsymbol{p}^*(\boldsymbol{s}) + \boldsymbol{y}_i^*(\boldsymbol{s}) \cdot \boldsymbol{q}^*(\boldsymbol{s}) - \boldsymbol{e}_i \cdot \boldsymbol{p}^*(\boldsymbol{s}) = 0$. Summing across the buyers, we get $\boldsymbol{p}^*(\boldsymbol{s}) \cdot \left(\sum_{i\in[n]} \boldsymbol{x}_i^*(\boldsymbol{s}) - \sum_{i\in[n]} \boldsymbol{e}_i\right) + \boldsymbol{q}^*(\boldsymbol{s}) \cdot \left(\sum_{i\in[n]} \boldsymbol{y}_i^*(\boldsymbol{s})\right) = 0$ for any $\boldsymbol{s} \in \mathcal{S}$, which is the Walras' law.

Finally, we want to show that $(\boldsymbol{X}^*, \boldsymbol{Y}^*, \boldsymbol{p}^*, \boldsymbol{q}^*)$ is feasible. We first show that $\sum_{i\in[n]} \boldsymbol{x}_i^*(\boldsymbol{s}) - \sum_{i\in[n]} \boldsymbol{e}_i \leq \boldsymbol{0}_m$ for any $\boldsymbol{s} \in \mathcal{S}$. We proved that for any state $\boldsymbol{s} \in \mathcal{S}$, $r'_{n+1}(\boldsymbol{s}, \boldsymbol{X}^*(\boldsymbol{s}), \boldsymbol{Y}^*(\boldsymbol{s}), \boldsymbol{p}^*(\boldsymbol{s}), \boldsymbol{q}^*(\boldsymbol{s})) = \boldsymbol{p}^*(\boldsymbol{s}) \cdot \left(\sum_{i\in[n]} \boldsymbol{x}_i^*(\boldsymbol{s}) - \sum_{i\in[n]} \boldsymbol{e}_i\right) + \boldsymbol{q}^*(\boldsymbol{s}) \cdot \left(\sum_{i\in[n]} \boldsymbol{y}_i^*(\boldsymbol{s})\right) = 0$, which implies $v_{n+1}^{\boldsymbol{\pi}^*}(\boldsymbol{s}) = 0$. For any $j \in [m]$, consider a $\boldsymbol{p} : \mathcal{S} \to \mathcal{P}$ defined by $\boldsymbol{p}(\boldsymbol{s}) = \boldsymbol{j}_j$ for all $\boldsymbol{s} \in \mathcal{S}$ and a $\boldsymbol{q} : \boldsymbol{s} \to \mathcal{Q}$ defined by $\boldsymbol{q}(\boldsymbol{s}) = \boldsymbol{0}_l$ for all $\boldsymbol{s} \in \mathcal{S}$. Then, we know that

$$
\begin{aligned}
0 &= v_{n+1}^{\boldsymbol{\pi}^*} && (14.49)\\
&= q_{n+1}^{\boldsymbol{\pi}^*}(\boldsymbol{s}, \boldsymbol{X}^*(\boldsymbol{s}), \boldsymbol{Y}^*(\boldsymbol{s}), \boldsymbol{p}^*(\boldsymbol{s}), \boldsymbol{q}^*(\boldsymbol{s})) && (14.50)\\
&\geq q_{n+1}^{\boldsymbol{\pi}^*}(\boldsymbol{s}, \boldsymbol{X}^*(\boldsymbol{s}), \boldsymbol{Y}^*(\boldsymbol{s}), \boldsymbol{p}(\boldsymbol{s}), \boldsymbol{q}(\boldsymbol{s})) && (14.51)\\
&= r'_{n+1}(\boldsymbol{s}, \boldsymbol{X}^*(\boldsymbol{s}), \boldsymbol{Y}^*(\boldsymbol{s}), \boldsymbol{p}(\boldsymbol{s}), \boldsymbol{q}(\boldsymbol{s})) + \mathop{\mathbb{E}}_{S' \sim \rho(S' | \boldsymbol{s}, \boldsymbol{Y}^*(\boldsymbol{s}))}[\gamma v_i^{\boldsymbol{\pi}^*}(S')] && (14.52)\\
&= \boldsymbol{j}_j \cdot \left(\sum_{i\in[n]} \boldsymbol{x}_i^*(\boldsymbol{s}) - \sum_{i\in[n]} \boldsymbol{e}_i\right) && \forall j \in [m] \quad (14.53)\\
&= \sum_{i\in[n]} x_{ij}^*(\boldsymbol{s}) - \sum_{i\in[n]} e_{ij} && \forall j \in [m] \quad (14.54)
\end{aligned}
$$

Thus, we know that $\sum_{i\in[n]} \boldsymbol{x}_i^*(\boldsymbol{s}) - \sum_{i\in[n]} \boldsymbol{e}_i \leq \boldsymbol{0}_m$ for any $\boldsymbol{s} \in \mathcal{S}$. Finally, we show that $\sum_{i\in[n]} \boldsymbol{y}_i^*(\boldsymbol{s}) \leq \boldsymbol{0}_l$ for all $\boldsymbol{s} \in \mathcal{S}$. By way of contradiction, suppose that for some asset $k \in [l]$, and some state $\boldsymbol{s} \in \mathcal{S}$, $\sum_{i\in[n]} y_{ik}^*(\boldsymbol{s}) > 0$. Then, the auctioneer can increase its cumulative payoff by increasing $q_k^*(\boldsymbol{s})$, which contradicts the definition of a GMPE.

Therefore, we can conclude that $\boldsymbol{\pi}^* = (\boldsymbol{X}^*, \boldsymbol{Y}^*, \boldsymbol{p}^*, \boldsymbol{q}^*) : \mathcal{S} \to \mathcal{X} \times \mathcal{Y} \times \mathcal{P} \times \mathcal{Q}$ is a RRE of $\mathcal{I}$.

Finally, notice that the transition functions set in our game are all stochastically concave and as such give rise action-value functions which are concave in the actions each of player (Atakan, 2003a), and it is easy to verify that the game also satisfies all conditions that guarantee the existence of a GMPE (see Section 4 of (Atakan, 2003a) for detailed proofs). Hence, by Theorem 12.2.1 which guarantees the existence of GMPE in generalized Markov games, we can conclude that there exists an RRE $(\boldsymbol{X}^*, \boldsymbol{Y}^*, \boldsymbol{p}^*, \boldsymbol{q}^*)$ in any Radner economy $\mathcal{I}$.

14.1.3 Omitted Results and Proofs from Section 13.1.4

Theorem 13.1.2.

Consider a Radner economy $\mathcal{I}$ and the associated Radner Markov pseudo-games $\mathcal{M}$. Let $(\boldsymbol{\pi}, \boldsymbol{\rho}, \mathbb{R}^{\Omega}, \mathbb{R}^{\Sigma})$ be a parametrization scheme for $\mathcal{M}$ and suppose Assumptions 12.3.2, 12.3.3, and 13.1.1 hold. Then, the convergence results in Theorem 12.3.1 hold for $\mathcal{M}$.

Proof

This results follows readily from Theorem 13.1.1 as an application of Theorem 12.3.1.

14.2 Experiments

14.2.1 Neural Projection Method

The projection method (Judd, 1992), also known as the weighted residual methods, is a numerical technique often used to approximate solutions to complex economic models, particularly those involving dynamic programming and dynamic stochastic general equilibrium (DSGE) models. These models are common in macroeconomics and often don't have analytical solutions due to their non-linear, dynamic, and high-dimensional nature. The projection method helps approximate these solutions by projecting the problem into a more manageable, lower-dimensional space.

The main idea of the projection method is to express equilibrium of the dynamic economic model as a solution to a functional equation $D(f) = \mathbf{0}$, where $f : \mathcal{S} \to \mathbb{R}^m$ is a function that represent some unknown policy, $D : (\mathcal{S} \to \mathbb{R}^m) \to (\mathcal{S} \to \mathbb{R}^n)$, and $\mathbf{0}$ is the zero vector. Some classic examples of the operator D includes Euler equations and Bellman equations. A canonical project method consists of four steps: 1) Define a set of basis functions $\{\psi_i : \mathcal{S} \to \mathbb{R}^m\}_{i\in[n]}$ and approximate each each function $f \in \mathcal{F}$ through a linear combination of basis functions: $\hat{f}(\cdot;\boldsymbol{\theta}) = \sum_{i=1}^{n} \theta_i \psi_i(\cdot)$; 2) Define a residual equation as a functional equation evaluated at the approximation: $R(\cdot;\boldsymbol{\theta}) \doteq D(\hat{f}(\cdot;\boldsymbol{\theta}))$; 3) Choose some weight functions $\{w_i : \mathcal{S} \to \mathbb{R}\}_{i\in[p]}$ over the states and find $\boldsymbol{\theta}$ that solves $F(\boldsymbol{\theta}) \doteq \int_{\mathcal{S}} w_i(\boldsymbol{s}) R(\boldsymbol{s};\boldsymbol{\theta}) d\boldsymbol{s} = 0$ for all $i \in [p]$. This gets the residual "close" to zero in the weighted integral sense; 4) Simulate the optimal decision rule based on the chosen parameter $\boldsymbol{\theta}$ and basis functions.

Recently, the neural projection method was developed to extend the traditional projection method (Maliar et al., 2021; Azinovic et al., 2022; Sauzet, 2021). In the neural projection method, neural networks are used as the functional approximators for policy functions instead of traditional basis function approximations. In this section, we show how we can approximate generalized Markov Perfect Nash equilibrium of generalized Markov game, and consequently Recursive Radner Equilibrium of Radner economies, through the neural projection method.

Assumption 14.2.1.

Given a generalized Markov game $\mathcal{M}$, assume that 1. for any $i \in [n]$, $\boldsymbol{s} \in \mathcal{S}$, $\boldsymbol{a}_{-i} \in \mathcal{A}_{-i}$, $\mathcal{X}_i(\boldsymbol{s}, \boldsymbol{a}_{-i}) \doteq \{\boldsymbol{a}_i \in \mathcal{A}_i \mid g_{ic}(\boldsymbol{s}, \boldsymbol{a}_i, \boldsymbol{a}_{-i}) \geq 0 \text{ for all } c \in [l]\}$ for a collection of **constraint functions** $\{g_{ic} : \mathcal{S} \times \mathcal{A} \mid c \in [l]\}$, where $\boldsymbol{a}_i \mapsto g_{ic}(\boldsymbol{s}, \boldsymbol{a}_i, \boldsymbol{a}_{-i})$ is concave for every $c \in [l]$.

Theorem 14.2.1.

Let $\mathcal{M}$ be a generalized Markov game that satisfies Assumption 14.2.1. For any policy profile $\boldsymbol{\pi} \in \mathcal{F}^{\text{markov}}$, $\boldsymbol{\pi}$ is a MPGNE if and only if there exists Lagrange multiplier policy $\boldsymbol{\lambda} : \mathcal{S} \to \mathbb{R}_+^{n\times l}$ such that $(\boldsymbol{\pi}, \boldsymbol{\lambda})$ solves the following functional equation: for all $i \in [n]$,

$s \in \mathcal{S}$,

$$0 \in \partial_{\boldsymbol{a}_i} q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{\pi}_i(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s})) + \sum_{c \in [l]} \lambda_{i,c}(\boldsymbol{s}) \partial_{\boldsymbol{a}_i} g_{ic}(\boldsymbol{s}, \boldsymbol{\pi}_i(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s})) \tag{14.55}$$

$$\forall c \in [l], \quad 0 = \lambda_{ic}(\boldsymbol{s}) g_{ic}(\boldsymbol{s}, \boldsymbol{\pi}_i(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s})) \tag{14.56}$$

$$\forall c \in [l], \quad 0 \leq g_{ic}(\boldsymbol{s}, \boldsymbol{a}_i^*, \boldsymbol{\pi}_{-i}(\boldsymbol{s})) \tag{14.57}$$

$$\tag{14.58}$$

and for all $i \in [n]$, $\boldsymbol{s} \in \mathcal{S}$,

$$v_i^{\boldsymbol{\pi}}(\boldsymbol{s}) = q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{\pi}_i(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s})) \tag{14.59}$$

Proof

First, we know that a policy profile $\boldsymbol{\pi} \in \mathcal{F}^{\text{markov}}$ is a MPGNE if and only if it satisfies the following generalized Bellman Optimality equations, i.e., for all $i \in [n]$, $\boldsymbol{s} \in \mathcal{S}$,

$$v_i^{\boldsymbol{\pi}}(\boldsymbol{s}) = \max_{\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{s}, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))} r_i(\boldsymbol{s}, \boldsymbol{a}_i, \boldsymbol{\pi}_{-i}(\boldsymbol{s})) + \mathbb{E}_{\boldsymbol{s}' \sim \rho(\cdot | \boldsymbol{s}, \boldsymbol{a}_i, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))}[\gamma v_i^{\boldsymbol{\pi}}(\boldsymbol{s}')] \tag{14.60}$$

$$= \max_{\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{s}, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))} q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{a}_i, \boldsymbol{\pi}_{-i}(\boldsymbol{s})) \tag{14.61}$$

Then, since $\boldsymbol{a}_i \mapsto q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{a}_i, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))$ is concave over $\mathcal{X}_i(\boldsymbol{s}, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))$ by Assumption 12.2.1, the KKT conditions provides sufficient and necessary optimality conditions for the constrained maximization problem

$$\max_{\boldsymbol{a}_i \in \mathcal{X}_i(\boldsymbol{s}, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))} q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{a}_i, \boldsymbol{\pi}_{-i}(\boldsymbol{s})) \tag{14.62}$$

That is, $\boldsymbol{a}_i^* \in \mathcal{X}_i(\boldsymbol{s}, \boldsymbol{\pi}_{-i}(\boldsymbol{s}))$ is a solution to eq. (14.62) if and only if there exists $\{\lambda_{ic}^* : \mathcal{S} \to \mathbb{R}_+\}_{c \in [l]}$ s.t.

$$0 \in \partial_{\boldsymbol{a}_i} q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{a}_i^*, \boldsymbol{\pi}_{-i}(\boldsymbol{s})) + \sum_{c \in [l]} \lambda_{ic}^*(\boldsymbol{s}) \partial_{\boldsymbol{a}_i} g_{ic}(\boldsymbol{s}, \boldsymbol{a}_i^*, \boldsymbol{\pi}_{-i}(\boldsymbol{s})) \tag{14.63}$$

$$\forall c \in [l], \quad 0 = \lambda_{ic}^*(\boldsymbol{s}) g_{ic}(\boldsymbol{s}, \boldsymbol{a}_i^*, \boldsymbol{\pi}_{-i}(\boldsymbol{s})) \tag{14.64}$$

$$\forall c \in [l], \quad 0 \leq g_{ic}(\boldsymbol{s}, \boldsymbol{a}_i^*, \boldsymbol{\pi}_{-i}(\boldsymbol{s})) \tag{14.65}$$

Therefore, we can conclude that $\boldsymbol{\pi} \in \mathcal{F}^{\text{markov}}$ is a MPGNE if and only if there exists $\{\lambda_{ic} : \mathcal{S} \to \mathbb{R}_+\}_{i\in[n],c\in[l]}$ s.t. for all $i \in [n]$, $\boldsymbol{s} \in \mathcal{S}$,

$$0 \in \partial_{\boldsymbol{a}_i} q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{\pi}_i(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s})) + \sum_{c\in[l]} \lambda_{i,c}(\boldsymbol{s}) \partial_{\boldsymbol{a}_i} g_{ic}(\boldsymbol{s}, \boldsymbol{\pi}_i(\boldsymbol{s}), \boldsymbol{\pi}_i(\boldsymbol{s})) \tag{14.66}$$

$$\forall c \in [l], \;\; 0 = \lambda_{ic}(\boldsymbol{s}) g_{ic}(\boldsymbol{s}, \boldsymbol{\pi}_i(\boldsymbol{s}), \boldsymbol{\pi}_i(\boldsymbol{s})) \tag{14.67}$$

$$\forall c \in [l], \;\; 0 \leq g_{ic}(\boldsymbol{s}, \boldsymbol{a}_i^*, \boldsymbol{\pi}_{-i}(\boldsymbol{s})) \tag{14.68}$$

and for all $i \in [n]$, $\boldsymbol{s} \in \mathcal{S}$,

$$v_i^{\boldsymbol{\pi}}(\boldsymbol{s}) = q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{\pi}_i(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s})) \tag{14.69}$$

Therefore, for a policy profile $\boldsymbol{\pi} \in \mathcal{F}^{\text{markov}}$ and a Lagrange multiplier policy $\boldsymbol{\lambda} : \mathcal{S} \to \mathbb{R}_+^{n\times l}$ such that $(\boldsymbol{\pi}, \boldsymbol{\lambda})$, consider the **total first order violation**

$$\Xi_{\text{first-order}}(\boldsymbol{\pi}, \boldsymbol{\lambda}) = \sum_{i\in[n]} \left\| \int_{\boldsymbol{s}\in\mathcal{S}} \partial_{\boldsymbol{a}_i} q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{\pi}_i(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s})) + \sum_{c\in[l]} \lambda_{i,c}(\boldsymbol{s}) \partial_{\boldsymbol{a}_i} g_{ic}(\boldsymbol{s}, \boldsymbol{\pi}_i(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s})) ds \right\|_2^2 \tag{14.70}$$

and the **average Bellman error**

$$\Xi_{\text{Bellman}}(\boldsymbol{\pi}, \boldsymbol{\lambda}) = \sum_{i\in[n]} \left\| \int_{\boldsymbol{s}\in\mathcal{S}} v_i^{\boldsymbol{\pi}}(\boldsymbol{s}) - q_i^{\boldsymbol{\pi}}(\boldsymbol{s}, \boldsymbol{\pi}_i(\boldsymbol{s}), \boldsymbol{\pi}_{-i}(\boldsymbol{s})) ds \right\|_2^2. \tag{14.71}$$

We can directly approximate the MPGNE through minimizing the sum of these two errors.

Typically, approximating the MPGNE using the neural projection method requires optimizing both the policy profile and the Lagrange multiplier policy. However, in exchange economy Markov pseudo-games, we derive a closed-form solution for the optimal Lagrange multiplier, allowing us to focus solely on optimizing the policy profile.

14.2.2 Implementation Details

Deterministic Case Training Details For deterministic transition probability case, for each reward function class we randomly sampled one economy with 10 consumers, 10 commodities, 1 asset, and 5 world state. The asset return matrix $\boldsymbol{R}$ is sampled in a way such that $r_{okj} \sim \text{Unif}([0.5, 1.1])$ for all o, k, and j. Moreover, we set the length of the stochastic process to be 30. For the initial state, we sample each consumer's endowment

$\boldsymbol{e}_i \sim \mathrm{Unif}([0.01, 0.1])^m$ and normalized so that the total endowment of each commodity add up to 1. We also sample each consumer's type $\boldsymbol{\theta}_i \sim \mathrm{Unif}([1.0, 5.0])^m$, and set the world state to be 0. The transition probability function ρ is defined as $\rho(\boldsymbol{s}' \mid \boldsymbol{s}, \boldsymbol{Y}) = 1$ for all $\boldsymbol{s}(o, \boldsymbol{E}, \boldsymbol{\Theta})$ where $\boldsymbol{s}' = (o', \boldsymbol{E}', \boldsymbol{\Theta}')$ is defined as $o' = 0$, $\boldsymbol{E}' = 0.01 \cdot \mathbf{1}_{n \times m}$, and $\boldsymbol{\Theta}' = \boldsymbol{\Theta}$.

Then, for both GAPNets method and neural projection method, we run 1000 episodes for each learning rate candidate in a grid search manner and measure the performance in terms of minimizing total first-order violation and average Bellman error. Finally, we pick the best hyperparameter for the final experiments.

In the final experiments, we run GAPNets for 2000 episodes using learning rates $\boldsymbol{\eta_\omega} = 1 \times 10^{-5}, \boldsymbol{\eta_\sigma} = 1 \times 10^{-5}$ for the linear economy, $\boldsymbol{\eta_\omega} = 1 \times 10^{-5}, \boldsymbol{\eta_\sigma} = 1 \times 10^{-5}$ for the Cobb-Douglas economy, and $\boldsymbol{\eta_\omega} = 1 \times 10^{-5}, \boldsymbol{\eta_\sigma} = 1 \times 10^{-5}$ for the Leontief economy. Similarly, we ran neural projection method for 2000 episodes using learning rates $\boldsymbol{\eta_\omega} = 1 \times 10^{-4}$ for the linear economy, $\boldsymbol{\eta_\omega} = 2.5 \times 10^{-5}$ for the Cobb-Douglas economy, and $\boldsymbol{\eta_\omega} = 1 \times 10^{-4}$ for the Leontief economy. In this process, we compute the exploitability of computed policy profile through gradient ascent of the adversarial network. In specific, we ran 1000 episodes of gradient ascent with learning rate $\boldsymbol{\eta_\sigma} = 5 \times 10^{-5}$ for the linear economy, $\boldsymbol{\eta_\sigma} = 1 \times 10^{-4}$ for the Cobb-Douglas economy, and $\boldsymbol{\eta_\sigma} = 1 \times 10^{-4}$ for the Leontief economy.

Next, for each economy, we randomly sample 50 policy profiles and record their total first-order violations, average Bellman errors, and exploitabilities. Finally, we normalize the results by the average of the sampled values.

Stochastic Case Training Details For stochastic transition probability case, for each reward function class we randomly sampled one economy with 10 consumers, 10 commodities, 1 asset, and 5 world state. The asset return matrix $\boldsymbol{R}$ is sampled in a way such that $r_{okj} \sim \mathrm{Unif}([0.5, 1.1])$ for all o, k, and j. Moreover, we set the length of the stochastic process to be 30. For the initial state, we sample each consumer's endowment $\boldsymbol{e}_i \sim \mathrm{Unif}([0.01, 0.1])^m$ and normalized so that the total endowment of each commodity add up to 1. We also sample each consumer's type $\boldsymbol{\theta}_i \sim \mathrm{Unif}([1.0, 5.0])^m$, and set the world state to be 0. The transition probability function will stochastically transition from state $\boldsymbol{s}(o, \boldsymbol{E}, \boldsymbol{\Theta})$ to state $\boldsymbol{s}' = (o', \boldsymbol{E}', \boldsymbol{\Theta}')$ where $o' \sim \mathrm{Unif}(\{0, 1, 2, 3, 4\})$, $\boldsymbol{E}' \sim 0.002 + \mathrm{Unif}([0.01, 0.1])^{n \times m}$, and $\boldsymbol{\Theta}' = \boldsymbol{\Theta}$.

Then, for both GAPNets method and neural projection method, we run 1000 episodes for each learning rate candidate in a grid search manner and measure the performance in terms of minimizing total first-order violation and average Bellman error. Finally, we pick the best hyperparameter for the final experiments.

In the final experiments, we run GAPNets for 2000 episodes using learning rates $\boldsymbol{\eta}_{\boldsymbol{\omega}} = 1 \times 10^{-5}, \boldsymbol{\eta}_{\boldsymbol{\sigma}} = 1 \times 10^{-5}$ for the linear economy, $\boldsymbol{\eta}_{\boldsymbol{\omega}} = 2.5 \times 10^{-5}, \boldsymbol{\eta}_{\boldsymbol{\sigma}} = 2.5 \times 10^{-5}$ for the Cobb-Douglas economy, and $\boldsymbol{\eta}_{\boldsymbol{\omega}} = 5 \times 10^{-5}, \boldsymbol{\eta}_{\boldsymbol{\sigma}} = 5 \times 10^{-5}$ for the Leontief economy. Similarly, we ran neural projection method for 2000 episodes using learning rates $\boldsymbol{\eta}_{\boldsymbol{\omega}} = 5 \times 10^{-5}$ for the linear economy, $\boldsymbol{\eta}_{\boldsymbol{\omega}} = 2.5 \times 10^{-5}$ for the Cobb-Douglas economy, and $\boldsymbol{\eta}_{\boldsymbol{\omega}} = 5 \times 10^{-4}$ for the Leontief economy. In this process, we compute the exploitability of computed policy profile through gradient ascent of the adversarial network. In specific, we ran 1000 episodes of gradient ascent with learning rate $\boldsymbol{\eta}_{\boldsymbol{\sigma}} = 7.5 \times 10^{-4}$ for the linear economy, $\boldsymbol{\eta}_{\boldsymbol{\sigma}} = 1 \times 10^{-4}$ for the Cobb-Douglas economy, and $\boldsymbol{\eta}_{\boldsymbol{\sigma}} = 1 \times 10^{-4}$ for the Leontief economy. When estimating the neural loss function—cumulative regret for the GAPNets method and total first-order violations and average Bellman error for the neural projection method—we use 100 samples for GAPNets and 10 samples for the neural projection method. The primary reason for this difference is the high memory consumption of the neural projection method, which makes larger sample sizes infeasible.

Next, for each economy, we randomly sample 50 policy profiles and record their total first-order violations, average Bellman errors, and exploitabilities. Finally, we normalize the results by the average of the sampled values.

14.2.3 Other Details

Programming Languages, Packages, and Licensing We ran our experiments in Python 3.7 (Van Rossum and Drake Jr, 1995), using NumPy (Harris et al., 2020), , CVXPY (Diamond and Boyd, 2016), Jax (Bradbury et al., 2018), OPTAX (Bradbury et al., 2018), Haiku (Hennigan et al., 2020), and JaxOPT (Blondel et al., 2021). All figures were graphed using Matplotlib (Hunter, 2007).

Python software and documentation are licensed under the PSF License Agreement. Numpy is distributed under a liberal BSD license. Pandas is distributed under a new BSD license. Matplotlib only uses BSD compatible code, and its license is based on the PSF license. CVXPY is licensed under an APACHE license.

Computational Resources The experiments were conducted using Google Colab, which provides cloud-based computational resources. Specifically, we utilized an NVIDIA T4 GPU with the following specifications: GPU: NVIDIA T4 (16GB GDDR6), CPU: Intel Xeon (2 vCPUs), RAM: 12GB, Storage: Colab-provided ephemeral storage.

Code Repository the full details of our experiments, including hyperparameter search, final experiment configurations, and visualization code, can be found in our code repository (`https://anonymous.4open.science/r/Markov-Pseudo-Game-EC2025-DCB8`).

Chapter 15

Conclusion

This thesis addresses the computational challenge of solving general equilibrium models by integrating techniques from computer science, optimization, and game theory. For over half a century, researchers have sought a general numerical method to compute equilibria in complex economies, a pursuit that began with Herbert Scarf's foundational work (Scarf, 1960), which itself was inspired by the works of Walras (1896), Uzawa (1960), and Arrow and Debreu (1954). While prior methods achieved partial success in small-scale models, no comprehensive solution existed for large, realistic economic systems. This thesis advances the field by introducing novel optimization frameworks, theoretical guarantees, and practical algorithms, offering a systematic resolution to this long-standing problem.

The first part of the thesis introduces variational inequalities (VIs) as a mathematical foundation for modeling Walrasian economies. A new class of first-order methods, the mirror extragradient algorithm, is developed, achieving polynomial-time convergence under the Minty condition and Bregman continuity. These results also include the first local convergence guarantees for constrained Bregman-continuous VIs without the Minty condition. This theoretical groundwork enables the formulation of the mirror *extratâtonnement* process, a price-adjustment mechanism that converges in all balanced Walrasian economies. This process provides the first polynomial-time, globally convergent price-adjustment method for computing Walrasian equilibria, resolving a major computational challenge in general equilibrium theory. Additionally, this thesis analyzes the convergence of *tâtonnement* in homothetic Fisher markets, offering a unified understanding of *tâtonnement* behavior across different utility functions.

Building on these results, the second part of the thesis extends the analysis to pseudo-games, a generalization of multiagent optimization frameworks, and their application to Arrow-Debreu economies. A new family of uncoupled learning dynamics, called mirror extragradient learning, is introduced, providing polynomial-time convergence guarantees for variationally stable pseudo-games. This part then reframes Arrow-Debreu equilibria as solutions to pseudo-games, enabling a computationally efficient characterization of equilibria in pure exchange economies through the application of mirror extragradient learning in the variationally stable trading post pseudo-game. This marks a significant advance in the computation of Arrow-Debreu equilibrium in pure exchange economies, for which no globally convergent market dynamics were known.

The final part of this thesis explores Markov pseudo-games, extending the previous frameworks to Radner economies, which explicitly incorporate time and uncertainty. This part extends pure equilibrium existence results in Markov games to settings with continuous action spaces, where previously only mixed-strategy equilibria were known. A novel learning-based approach is then introduced for computing generalized Markov perfect equilibria (GMPE), leveraging adversarial learning techniques to compute equilibrium policies in polynomial time. Focus then shifts to computing Radner equilibrium, an inherently infinite-dimensional problem. A function-approximation method inspired by merit functions is introduced, allowing for efficient computation of a solution that satisfies the necessary conditions of a Radner equilibrium, under suitable smoothness conditions. These findings open a new research direction at the intersection of deep learning, reinforcement learning, and mathematical economics, offering a promising path toward scalable, data-driven economic models.

The ordering of the three major parts is intentional and serves three key purposes. First, the results on VIs form the mathematical backbone of the thesis and are used throughout, requiring them to be presented first. Second, the fact that Walrasian economies can be seen as a special case of Arrow-Debreu economies suggests a natural progression from one to the other, with insights from Walrasian models helping to contextualize results in Arrow-Debreu markets. Finally, Markov pseudo-games and Radner economies extend the previous models to infinite-dimensional settings, where equilibrium solutions become significantly more complex and require modern learning-based computational approaches. This natural progression not only unifies classical and modern equilibrium models, but also provides a clear research trajectory toward infinite-dimensional optimization and economies with infinitely many commodities, a rapidly growing field in economic theory and applied mathematics.

In summary, this thesis provides a unified computational framework for solving general equilibrium models, integrating classical economic theory with modern optimization and learning techniques. It resolves long-standing computational challenges in computing Walrasian, Arrow-Debreu, and Radner equilibria, while also introducing novel algorithmic tools that have broader implications for optimization, game theory, and artificial intelligence.

15.1 Future Directions

I now highlight several promising directions for future research that I find both exciting and relevant to the public good. Rather than following the chronological order in which they were presented in this thesis, I have organized them based on their potential to advance research and foster the application of general equilibrium theory in real-world settings.

15.1.1 Radner Economies and Infinite-Dimensional Walrasian Economies

The Walrasian and Arrow-Debreu economy models studied in this thesis were finite-dimensional in that they considered only a finite set of commodities. More recently, infinite-dimensional generalizations of these economies have been explored (Prescott and Lucas, 1972), and equilibrium existence has been established. However, apart from one notable work (Gao and Kroer, 2021), little is known about computing Walrasian equilibria in infinite-dimensional settings. Advancing this research is crucial for applying general equilibrium models to policy analysis, as macroeconomic policy models often take the form of Radner economies, which themselves can be viewed as special cases of infinite-dimensional Walrasian economies. Specifically, any Radner economy can be reformulated as a Walrasian economy in which the set of goods is given by the union of all commodities and assets across all states. Since the state space in Radner economies is typically continuous, the resulting Walrasian economy is generally infinite-dimensional. Indeed, the purpose of the study of Radner economies in this thesis was to push the envelope of algorithmic general equilibrium towards infinite-dimensional economies.

In Part III, I propose one approach to solving infinite-dimensional economies, but this merely scratches the surface of what is possible. Given the complexity of these problems, machine learning and artificial intelligence are likely to play a pivotal role in their resolution. Future research should focus on developing methods for even more complex stochastic and high-dimensional economies, incorporating deep learning techniques to improve scalability and robustness. This emerging field, at the intersection of deep learning, reinforcement learning, optimization theory, and mathematical economics, holds immense potential. More importantly, it offers a long-overdue opportunity to bring the transformative power of AI to economic policy-making—an area that has remained largely untouched by these advances for far too long.

15.1.2 Walrasian Economies

The theoretical insights and empirical validation presented in Part I suggest that the computational intractability of general equilibrium problems arises primarily from discontinuities rather than inherent complexity. This perspective challenges long-standing assumptions in applied general equilibrium theory, particularly those stemming from the works of Scarf (1960), Papadimitriou and Yannakakis (2010), Codenotti et al. (2006), and Deng and Du (2008). Our empirical results indicate that the mirror *extratâtonnement* process *can* compute Walrasian equilibria in cases where prior theoretical results suggested this would be infeasible (e.g., Leontief Arrow-Debreu economies (Codenotti et al., 2006; Deng and Du, 2008)). While I have attempted to explain these observations through the pathwise Bregman-smoothness assumption, future work should explore specific kernel functions and characterize the class of Walrasian economies where this assumption holds, as past work has done (Goktas et al., 2023c). Doing so would provide deeper insights into the computational challenges identified over the past two decades.

Moreover, our experiments demonstrate the ability of the mirror *extratâtonnement* process to solve very large Leontief Arrow-Debreu economies. Since these economies are known to be PPAD-complete (Deng and Du, 2008; Codenotti et al., 2006), there exists a polynomial-time reduction from games to Leontief Arrow-Debreu economies. This suggests that the mirror *extratâtonnement* process could also be used to solve large games in practice. Future research should investigate such algorithms, potentially exploring formulations of Nash equilibria as discontinuous variational inequalities satisfying the Minty condition.

15.1.3 Arrow-Debreu Economies

While the trading post pseudo-game provides a tractable characterization of Arrow-Debreu equilibria in pure exchange economies, it remains an open question whether this approach extends to concave Arrow-Debreu economies. Future work should explore whether the trading post pseudo-game can be generalized to characterize all Arrow-Debreu economies, thereby paving the way for market dynamics that converge to equilibrium in a broader class of Arrow-Debreu economies.

15.1.4 Fisher Markets

Future work could investigate the space of homogeneous utility functions with negative cross-price elasticity of Hicksian demand to possibly derive faster convergence rates than those provided in this paper. Additionally, it remains to be seen if the bound we have provided in this paper is tight; the greatest lower bound known for the convergence of *tâtonnement* in homothetic Fisher markets is $O(1/t^2)O(1/\varepsilon^2)$ for Leontief markets (Cheung et al., 2012), leaving space for improvement. Finally, Lemma 6.4.1 (Chapter 6) suggests that to extend convergence results for *tâtonnement* beyond homothetic domains, one might have to consider the Hicksian demand elasticity *w.r.t. utility level* rather than price.

15.1.5 Variational Inequalities

The polynomial-time best-iterate convergence of the mirror extragradient method to an approximate strong solution has been established under the assumption that the kernel function is strongly convex and Lipschitz-smooth. However, it is likely that the Lipschitz-smoothness assumption can be dropped, as it does not appear necessary to show that the distance between intermediate and final iterates decreases. This generalization would be highly useful, given that many Bregman divergences are generated by strongly convex functions that are not Lipschitz-smooth (e.g., Itakura-Saito Divergence (Itakura and Saito, 1968)). Another promising direction is to investigate whether the mirror extragradient method can be shown to converge in polynomial time to an approximate strong solution in last iterates under the Minty condition.

Bibliography

Sydney N Afriat. The construction of utility functions from expenditure data. *International economic review*, 8(1):67–77, 1967. (Cited on pages 66 and 146.)

Sydney. N. Afriat. Theory of maxima and the method of lagrange. *SIAM Journal on Applied Mathematics*, 20(3):343–357, 1971. ISSN 00361399. URL http://www.jstor.org/stable/2099955. (Cited on page 178.)

Alekh Agarwal, Sham M Kakade, Jason D Lee, and Gaurav Mahajan. Optimality and approximation with policy gradient methods in markov decision processes. In *Conference on Learning Theory*, pages 64–66. PMLR, 2020. (Cited on pages 315 and 318.)

Gabriel Andrade, Rafael Frongillo, Sharadha Srinivasan, and Elliot Gorokhovsky. Graphical economies with resale. *Proceedings of the 22nd ACM Conference on Economics and Computation*, Jul 2021. doi: 10.1145/3465456.3467628. URL http://dx.doi.org/10.1145/3465456.3467628. (Cited on page 131.)

Danilo Ardagna, Barbara Panicucci, and Mauro Passacantando. A game theoretic formulation of the service provisioning problem in cloud systems. In *Proceedings of the 20th International Conference on World Wide Web*, WWW '11, page 177–186, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450306324. doi: 10.1145/1963405.1963433. URL https://doi.org/10.1145/1963405.1963433. (Cited on page 218.)

Danilo Ardagna, Michele Ciavotta, and Mauro Passacantando. Generalized nash equilibria for the service provisioning problem in multi-cloud systems. *IEEE Transactions on Services*

Computing, 10(3):381–395, 2017. doi: 10.1109/TSC.2015.2477836. (Cited on page 218.)

Kazeem Olalekan Aremu, Muhammed Ibrahim Mona, and Muhammad Ibrahim. A modified self-adaptive inertial tseng algorithm for solving a quasimonotone variational inequality and fixed point problems in real hilbert space. *The Journal of Analysis*, pages 1–22, 2024. (Cited on page 84.)

Kenneth Arrow. An extension of the basic theorems of classical welfare economics. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*. The Regents of the University of California, 1951a. (Cited on pages 21 and 219.)

Kenneth Arrow and Gerard Debreu. Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society*, pages 265–290, 1954. (Cited on pages 1, 3, 17, 18, 19, 20, 65, 69, 115, 119, 126, 155, 217, 218, 219, 222, 225, 232, 262, 264, 266, 294, 295, 299, 302, 306, and 353.)

Kenneth J Arrow. An extension of the basic theorems of classical welfare economics. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, volume 2, pages 507–533. University of California Press, 1951b. (Cited on pages 219 and 296.)

Kenneth J Arrow. Le role des valeurs boursieres pour la repartition la meilleure des risques, econometrie, 41-47, english translation as the role of securities in the optimal allocation of risk-bearing. *Review of Economic Studies*, 31:91–96, 1964. (Cited on pages 19, 295, 296, 297, and 298.)

Kenneth J. Arrow and Leonid Hurwicz. On the stability of the competitive equilibrium, i. *Econometrica*, 26(4):522–552, 1958. ISSN 00129682, 14680262. URL `http://www.jstor.org/stable/1907515`. (Cited on pages 65, 66, 83, 119, 132, 133, 146, 150, and 219.)

Kenneth J Arrow and Leonid Hurwicz. Competitive stability under weak gross substitutability: The" euclidean distance" approach. *International Economic Review*, 1(1):38–49, 1960. (Cited on page 66.)

Kenneth J. Arrow and Timothy J. Kehoe. Distinguished fellow: Herbert scarf's contributions to economics. *The Journal of Economic Perspectives*, 8(4):161–181, 1994. ISSN 08953309. URL `http://www.jstor.org/stable/2138344`. (Cited on page 66.)

Kenneth J Arrow and Marc Nerlove. A note on expectations and stability. *Econometrica: Journal of the Econometric Society*, pages 297–305, 1958. (Cited on page 219.)

Kenneth J. Arrow, Leonid Hurwicz, and Hirofumi Uzawa. *Studies in Linear and Non-Linear Programming*, volume 2 of *Stanford Mathematical Studies in the Social Sciences*. Stanford University Press, Redwood City, California, 1958. (Cited on pages 83, 311, and 315.)

Kenneth J Arrow, Henry D Block, and Leonid Hurwicz. On the stability of the competitive equilibrium, ii. *Econometrica: Journal of the Econometric Society*, pages 82–109, 1959. (Cited on pages 66 and 146.)

Kenneth Joseph Arrow, Frank Hahn, et al. General competitive analysis. 1971. (Cited on pages 164 and 165.)

Alp E Atakan. Stochastic convexity in dynamic programming. *Economic Theory*, 22:447–455, 2003a. (Cited on pages 345 and 346.)

Alp E. Atakan. Stochastic convexity in dynamic programming. *Economic Theory*, 22(2):447–455, 2003b. ISSN 09382259, 14320479. URL `http://www.jstor.org/stable/25055693`. (Cited on page 302.)

Jean-Pierre Aubin. *Optima and equilibria: an introduction to nonlinear analysis*, volume 140. Springer Science & Business Media, 2013. (Cited on page 240.)

Adrien Auclert, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5): 2375–2408, 2021. (Cited on page 249.)

Alfred Auslender and Marc Teboulle. Interior projection-like methods for monotone variational inequalities. *Mathematical programming*, 104(1):39–68, 2005. (Cited on page 84.)

Didier Aussel and Parin Chaipunya. Variational and quasi-variational inequalities under local reproducibility: solution concept and applications. *Journal of Optimization Theory and Applications*, 203(2):1531–1563, 2024. (Cited on pages 96 and 97.)

Didier Aussel, Jean-Noël Corvellec, and Marc Lassonde. Subdifferential characterization of quasiconvexity and convexity. *J. Convex Anal*, 1(2):195–201, 1994. (Cited on pages 255 and 282.)

Marlon Azinovic, Luca Gaegauf, and Simon Scheidegger. Deep equilibrium nets. *International Economic Review*, 63(4):1471–1525, 2022. (Cited on pages 329 and 347.)

Qin Ba and Jong-Shi Pang. Exact penalization of generalized nash equilibrium problems. *Operations Research*, 2020. (Cited on pages 219 and 233.)

Yves Balasko. Some results on uniqueness and on stability of equilibrium in general equilibrium theory. *Journal of Mathematical Economics*, 2(2):95–118, 1975. (Cited on pages 65 and 219.)

Xiaohui Bei, Jugal Garg, and Martin Hoefer. Tatonnement for linear and gross substitutes markets. *CoRR abs/1507.04925*, 2015. (Cited on page 130.)

Richard Bellman. Dynamic programming. *Science*, 153(3731):34–37, 1966. (Cited on page 306.)

Aharon Ben-Tal, Laurent El Ghaoui, and Arkadi Nemirovski. *Robust optimization*, volume 28. Princeton university press, 2009. (Cited on page 75.)

Claude Berge. *Topological Spaces: including a treatment of multi-valued functions, vector spaces, and convexity*. Courier Corporation, 1997. (Cited on pages 58, 59, 184, 186, 211, 225, 226, 251, 252, 267, and 273.)

Dimitri P Bertsekas. Incremental gradient, subgradient, and proximal methods for convex optimization: A survey. 2011. (Cited on page 51.)

Jalaj Bhandari and Daniel Russo. Global optimality guarantees for policy gradient methods. *arXiv preprint arXiv:1906.01786*, 2019. (Cited on pages 307 and 342.)

Benjamin Birnbaum, Nikhil R. Devanur, and Lin Xiao. Distributed algorithms via gradient descent for Fisher markets. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, EC '11, page 127–136, New York, NY, USA, 2011. Association for Computing Machinery. ISBN 9781450302616. doi: 10.1145/1993574.1993594. URL `https://doi.org/10.1145/1993574.1993594`. (Cited on pages 163 and 164.)

Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal of political economy*, 81(3):637–654, 1973. (Cited on page 297.)

David Blackwell. Discounted dynamic programming. *The Annals of Mathematical Statistics*, 36(1):226–235, 1965. (Cited on page 323.)

Mathieu Blondel, Quentin Berthet, Marco Cuturi, Roy Frostig, Stephan Hoyer, Felipe Llinares-López, Fabian Pedregosa, and Jean-Philippe Vert. Efficient and modular implicit differentiation. *arXiv preprint arXiv:2105.15183*, 2021. (Cited on pages 191 and 352.)

Lawrence E. Blume. Duality. *The New Palgrave Dictionary of Economics*, pages 1–7, 2017. doi: 10.1057/978-1-349-95121-5_285-2. URL `https://doi.org/10.1057/978-1-349-95121-5_285-2`. (Cited on pages 170, 172, 197, and 199.)

KC Border. Euler's theorem for homogeneous functions. *Caltech Division of the Humanities and Social Sciences*, 27:16–34, 2017. (Cited on page 199.)

Kim Border. Introduction to correspondences, September 2010. (Cited on page 36.)

Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004. (Cited on pages 50, 57, 131, 163, 171, 279, and 282.)

James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL `http://github.com/google/jax`. (Cited on pages 191 and 352.)

William C Brainard, Herbert E Scarf, et al. *How to compute equilibrium prices in 1891*. Citeseer, 2000. (Cited on pages 128, 161, and 166.)

Simina Brânzei, Nikhil Devanur, and Yuval Rabani. Proportional dynamics in exchange economies. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 180–201, 2021. (Cited on page 131.)

Lev M Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR computational mathematics and mathematical physics*, 7(3):200–217, 1967. (Cited on page 163.)

Michele Breton, Georges Zaccour, and Mehdi Zahaf. A game-theoretic formulation of joint implementation of environmental projects. *European Journal of Operational Research*, 168 (1):221–239, 2006. (Cited on page 218.)

Haim Brezis and Haim Brezis. *Functional analysis, Sobolev spaces and partial differential equations*, volume 2. Springer, 2011. (Cited on page 83.)

Haïm Brezis and Moïse Sibony. Méthodes d'approximation et d'itération pour les opérateurs monotones. In *Numerical Analysis of Partial Differential Equations*, pages 407–439. Springer, 2011. (Cited on page 75.)

Felix E Browder. Nonlinear monotone operators and convex sets in banach spaces. *Bulletin of the American Mathematical Society*, 71(5):780–785, 1965. (Cited on page 75.)

Michael Bruckner and Tobias Scheffer. Nash equilibria of static prediction games. *Advances in neural information processing systems*, 22, 2009. (Cited on page 218.)

Michael Bruckner, Christian Kanzow, and Tobias Scheffer. Static prediction games for adversarial learning problems. *J. Mach. Learn. Res.*, 13:2617–2654, 2012. URL http://dl.acm.org/citation.cfm?id=2503326. (Cited on page 218.)

Luis Felipe Bueno, Gabriel Haeser, and Frank Navarro Rojas. Optimality conditions and constraint qualifications for generalized nash equilibrium problems and their practical implications. *SIAM Journal on Optimization*, 29(1):31–54, 2019. doi: 10.1137/17M1162524. (Cited on page 233.)

U.S. Department of Commerce Bureau of Economic Analysis. National income and product accounts tables, 1940–1947, 2021. URL http://bea.gov/iTable/iTable.cfm?ReqID=9&step=1. GDP Accounts in Billions of Chained 1937 Dollars. (Cited on page 17.)

J. C. Butcher. *Numerical Methods for Ordinary Differential Equations*. John Wiley & Sons, 2nd edition, 2008. ISBN 9780470753751. (Cited on page 156.)

Yang Cai, Argyris Oikonomou, and Weiqiang Zheng. Tight last-iterate convergence of the extragradient method for constrained monotone variational inequalities. *arXiv preprint arXiv:2204.09228*, 2022. (Cited on pages 82, 83, 84, 87, and 88.)

Richard Cantillon. Essai sur la nature du commerce en général. *Fetcher Gyler*, 1755. (Cited on page 15.)

Augustin Cauchy et al. Méthode générale pour la résolution des systemes d'équations simultanées. *Comp. Rend. Sci. Paris*, 25(1847):536–538, 1847. (Cited on pages 83 and 86.)

Augustin-Louis Cauchy. *Cours d'Analyse de l'École Royale Polytechnique*. Debure, Paris, 1821. In French. Contains the first formal appearance of what is now the Cauchy-Schwarz inequality. (Cited on page 90.)

Yair Censor, Aviv Gibali, and Simeon Reich. Strong convergence of subgradient extragradient methods for the variational inequality problem in hilbert space. *Optimization Methods and Software*, 26(4-5):827–845, 2011. (Cited on page 84.)

Xi Chen and Xiaotie Deng. 3-nash is ppad-complete. In *Electronic Colloquium on Computational Complexity*, volume 134, pages 2–29. Citeseer, 2005. (Cited on page 3.)

Xi Chen and Shang-Hua Teng. Spending is not easier than trading: on the computational equivalence of Fisher and Arrow-Debreu equilibria. In *International Symposium on Algorithms and Computation*, pages 647–656. Springer, 2009. (Cited on pages 5, 65, and 219.)

Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *Journal of the ACM (JACM)*, 56(3):1–57, 2009. (Cited on pages 144 and 309.)

Yun Kuen Cheung. *Analyzing tatonnement dynamics in economic markets*. PhD thesis, New York University, 2014. (Cited on pages 73 and 176.)

Yun Kuen Cheung and Richard Cole. Amortized analysis of asynchronous price dynamics. *arXiv preprint arXiv:1806.10952*, 2018. (Cited on page 130.)

Yun Kuen Cheung, Richard Cole, and Ashish Rastogi. Tatonnement in ongoing markets of complementary goods. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 337–354, 2012. (Cited on pages 130 and 358.)

Yun Kuen Cheung, Richard Cole, and Nikhil Devanur. Tatonnement beyond gross substitutes? gradient descent to the rescue. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*, STOC '13, page 191–200, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450320290. doi: 10.1145/2488608.2488633. URL `https://doi.org/10.1145/2488608.2488633`. (Cited on pages 65, 73, 130, 131, 138, 150, 164, 171, 174, 176, 181, 200, 202, and 219.)

Yun Kuen Cheung, Richard Cole, and Yixin Tao. Dynamics of distributed updating in Fisher markets. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 351–368, 2018. (Cited on pages 65, 131, 138, 163, and 219.)

Yun Kuen Cheung, Martin Hoefer, and Paresh Nakhe. Tracing equilibrium in dynamic markets via distributed adaptation. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 1225–1233, 2019. (Cited on page 131.)

Francis Clarke. Necessary conditions in optimal control and in the calculus of variations. In *Differential equations, chaos and variational problems*, pages 143–156. Springer, 2007. (Cited on pages 45 and 50.)

Frank H Clarke. *Optimization and nonsmooth analysis*. SIAM, 1990. (Cited on pages 44 and 45.)

Bruno Codenotti, Benton McCune, and Kasturi Varadarajan. Market equilibrium via the excess demand function. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 74–83, 2005. (Cited on pages 15, 65, 68, 72, 73, 129, 130, and 219.)

Bruno Codenotti, Amin Saberi, Kasturi Varadarajan, and Yinyu Ye. Leontief economies encode nonzero sum two-player games. In *SODA*, volume 6, pages 659–667, 2006. (Cited on pages 65, 70, 143, 219, and 357.)

Richard Cole and Lisa Fleischer. Fast-converging tatonnement algorithms for one-time and ongoing market problems. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 315–324, 2008. (Cited on pages 15, 65, 130, 139, 175, and 219.)

Richard Cole and Yixin Tao. Balancing the robustness and convergence of tatonnement, 2019. (Cited on pages 73 and 175.)

Richard Cole, Lisa Fleischer, and Ashish Rastogi. Discrete price updates yield fast convergence in ongoing markets with finite warehouses, 2010. (Cited on page 130.)

S Cook. The complexity of theorem proving procedures. In *Proc. 3rd Annual ACM Symposium on the Theory of Computation*, pages 151–158, 1971. (Cited on page 2.)

Richard W. Cottle and George B. Dantzig. Complementary pivot theory of mathematical programming. *Linear Algebra and Its Applications*, 1(1):103–125, 1968. (Cited on page 117.)

Eleftherios Couzoudis and Philipp Renner. Computing generalized nash equilibria by polynomial programming. *Mathematical Methods of Operations Research*, 77(3):459–472, 2013. (Cited on page 219.)

Giovanni P Crespi, Angelo Guerraggio, and Matteo Rocca. Minty variational inequality and optimization: scalar and vector case. In *Generalized Convexity, Generalized Monotonicity and Applications: Proceedings of the 7 th International Symposium on Generalized Convexity and Generalized Monotonicity*, pages 193–211. Springer, 2005. (Cited on pages 77 and 79.)

Stella Dafermos. Exchange price equilibria and variational inequalities. *Mathematical Programming*, 46(1):391–402, 1990. (Cited on pages 1, 71, and 117.)

John M. Danskin. The theory of max-min, with applications. *SIAM Journal on Applied Mathematics*, 14(4):641–664, 1966. ISSN 00361399. URL `http://www.jstor.org/stable/2946123`. (Cited on pages 106, 172, 183, and 196.)

Constantinos Daskalakis. Non-concave games: A challenge for game theory's next 100 years. *Cowles Preprints*, 2022. (Cited on pages 221 and 257.)

Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009. (Cited on pages 3 and 309.)

Constantinos Daskalakis, Dylan J Foster, and Noah Golowich. Independent policy gradient methods for competitive reinforcement learning. *Advances in neural information processing systems*, 33:5527–5540, 2020a. (Cited on pages 300, 315, 316, 341, and 342.)

Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. The complexity of constrained min-max optimization, 2020b. (Cited on pages 48 and 257.)

Damek Davis and Dmitriy Drusvyatskiy. Stochastic model-based minimization of weakly convex functions. *SIAM Journal on Optimization*, 29(1):207–239, 2019. (Cited on pages 50 and 51.)

Damek Davis, Dmitriy Drusvyatskiy, Kellie J MacPhee, and Courtney Paquette. Subgradient methods for sharp weakly convex functions. *Journal of Optimization Theory and Applications*, 179:962–982, 2018. (Cited on pages 50 and 341.)

Gerard Debreu. The coefficient of resource utilization. *Econometrica*, 19(3):273–292, 1951a. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1906814. (Cited on pages 219 and 296.)

Gerard Debreu. The coefficient of resource utilization. *Econometrica: Journal of the Econometric Society*, pages 273–292, 1951b. (Cited on page 21.)

Gerard Debreu. Excess demand functions. *Journal of mathematical economics*, 1(1):15–21, 1974. (Cited on pages 119 and 123.)

Gerard Debreu et al. Representation of a preference ordering by a numerical function. *Decision processes*, 3:159–165, 1954. (Cited on pages 262 and 320.)

Xiaotie Deng and Ye Du. The computation of approximate competitive equilibrium is ppad-hard. *Information Processing Letters*, 108(6):369–373, 2008. (Cited on pages 5, 70, 144, and 357.)

N. R. Devanur, C. H. Papadimitriou, A. Saberi, and V. V. Vazirani. Market equilibrium via a primal-dual-type algorithm. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 389–395, 2002. doi: 10.1109/SFCS.2002.1181963. (Cited on pages 128 and 168.)

Nikhil R Devanur, Christos H Papadimitriou, Amin Saberi, and Vijay V Vazirani. Market equilibrium via a primal–dual algorithm for a convex program. *Journal of the ACM (JACM)*, 55(5):1–18, 2008. (Cited on pages 128, 131, and 171.)

Nikhil R Devanur, Kamal Jain, Tung Mai, Vijay V Vazirani, and Sadra Yazdanbod. New convex programs for fisher's market model and its generalizations. *arXiv preprint arXiv:1603.01257*, 2016. (Cited on page 170.)

Jelena Diakonikolas. Halpern iteration for near-optimal and parameter-free monotone inclusion and strong solutions to variational inequalities. In *Conference on Learning Theory*, pages 1428–1451. PMLR, 2020. (Cited on pages 76, 83, and 84.)

Peter A. Diamond. The role of a stock market in a general equilibrium model with technological uncertainty. *The American Economic Review*, 57(4):759–776, 1967. ISSN 00028282. URL http://www.jstor.org/stable/1815367. (Cited on page 297.)

Steven Diamond and Stephen Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016. (Cited on page 352.)

Egbert Dierker. Chapter 17 regular economies. volume 2 of *Handbook of Mathematical Economics*, pages 795–830. Elsevier, 1982. doi: https://doi.org/10.1016/S1573-4382(82)02012-8. URL https://www.sciencedirect.com/science/article/pii/S1573438282020128. (Cited on page 219.)

Jean Dieudonné. *Treatise on Analysis: Volume 1*, volume 1. Academic press, 1960. (Cited on page 32.)

Axel Dreves. Computing all solutions of linear generalized nash equilibrium problems. *Mathematical Methods of Operations Research*, 85(2):207–221, 2017. (Cited on pages 219 and 232.)

Axel Dreves and Nathan Sudermann-Merx. Solving linear generalized nash equilibrium problems numerically. *Optimization Methods and Software*, 31(5):1036–1063, 2016. (Cited on page 219.)

Axel Dreves, Anna Heusinger, Christian Kanzow, and Masao Fukushima. A globalized newton method for the computation of normalized nash equilibria. *J. of Global Optimization*, 56(2):327–340, 6 2013. ISSN 0925-5001. doi: 10.1007/s10898-011-9824-9. URL `https://doi.org/10.1007/s10898-011-9824-9`. (Cited on page 232.)

Jacques H Dreze. Investment under private ownership: optimality, equilibrium and stability. In *Allocation under Uncertainty: Equilibrium and Optimality: Proceedings from a Workshop sponsored by the International Economic Association*, pages 129–166. Springer, 1974. (Cited on page 298.)

Dmitriy Drusvyatskiy and Courtney Paquette. Efficiency of minimizing compositions of convex functions and smooth maps. *Mathematical Programming*, 178:503–558, 2019. (Cited on page 108.)

Vu Tien Dung, Pham Ky Anh, and Duong Viet Thong. Convergence of two-step inertial tseng's extragradient methods for quasimonotone variational inequality problem. *Communications in Nonlinear Science and Numerical Simulation*, page 108110, 2024. (Cited on page 84.)

G. Duvaut and J.-L. Lions. *Inequalities in Mechanics and Physics*, volume 219 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1976. doi: 10.1007/978-3-642-66282-9. (Cited on page 1.)

Edmund Eisenberg. Aggregation of utility functions. *Management Science*, 7(4):337–350, 1961. (Cited on page 128.)

Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959. (Cited on page 168.)

Francisco Facchinei and Christian Kanzow. Generalized nash equilibrium problems. *Annals of Operations Research*, 175(1):177–211, 2010a. (Cited on pages 218, 219, 226, and 299.)

Francisco Facchinei and Christian Kanzow. Penalty methods for the solution of generalized nash equilibrium problems. *SIAM Journal on Optimization*, 20(5):2228–2253, 2010b. doi: 10.1137/090749499. URL https://doi.org/10.1137/090749499. (Cited on page 233.)

Francisco Facchinei and Lorenzo Lampariello. Partial penalization for the solution of generalized nash equilibrium problems. *Journal of Global Optimization*, 50(1):39–57, 2011. (Cited on pages 219 and 232.)

Francisco Facchinei and Jong-Shi Pang. *Finite-dimensional variational inequalities and complementarity problems*. Springer, 2003. (Cited on pages 75, 78, and 81.)

Francisco Facchinei and Simone Sagratella. On the computation of all solutions of jointly convex generalized nash equilibrium problems. *Optimization Letters*, 5(3):531–547, 2011. (Cited on page 219.)

Francisco Facchinei, Andreas Fischer, and Veronica Piccialli. On generalized nash games and variational inequalities. *Operations Research Letters*, 35(2):159–164, 2007. ISSN 0167-6377. doi: https://doi.org/10.1016/j.orl.2006.03.004. URL https://www.sciencedirect.com/science/article/pii/S0167637706000484. (Cited on page 232.)

Francisco Facchinei, Andreas Fischer, and Veronica Piccialli. Generalized nash equilibrium problems and newton methods. *Mathematical Programming*, 117(1):163–194, 2009. (Cited on pages 219 and 232.)

Jesús Fernández-Villaverde. *Computational Methods for Macroeconomics*. University of Pennsylvania, 2023. URL https://www.sas.upenn.edu/~jesusfv/teaching.html. Lecture notes for one-year course on computational methods for economists. (Cited on pages 68, 299, and 329.)

Anthony V Fiacco and Jerzy Kyparisis. Convexity and concavity properties of the optimal value function in parametric nonlinear programming. *Journal of optimization theory and applications*, 48(1):95–126, 1986. (Cited on page 333.)

Arlington M Fink. Equilibrium in a stochastic n-person game. *Journal of science of the hiroshima university, series ai (mathematics)*, 28(1):89–93, 1964. (Cited on pages 11, 299, and 302.)

Andreas Fischer, Markus Herrich, Alexey F Izmailov, and Mikhail V Solodov. A globally convergent lp-newton method. *SIAM Journal on Optimization*, 26(4):2012–2033, 2016. (Cited on pages 219 and 232.)

Franklin M Fisher. The stability of general equilibrium: Results and problems. *1976), Essays in Economic Analysis: Proceedings of the Association of University Teachers of Economics Sheffield*, pages 3–29, 1975. (Cited on page 67.)

Sjur Flam and Andrzej Ruszczynski. Noncooperative convex games: Computing equilibrium by partial regularization. Working papers, International Institute for Applied Systems Analysis, 1994. URL https://EconPapers.repec.org/RePEc:wop:iasawp:wp94042. (Cited on pages 241, 244, and 309.)

Duncan K Foley. Lindahl's solution and the core of an economy with public goods. *Econometrica: Journal of the Econometric Society*, pages 66–72, 1970. (Cited on page 22.)

Gerald B Folland. *Real analysis: modern techniques and their applications*, volume 40. John Wiley & Sons, 1999. (Cited on pages 27 and 39.)

Lance Fortnow. The golden ticket: P, np, and the search for the impossible. In *The Golden Ticket*. Princeton University Press, 2013. (Cited on page 2.)

Masao Fukushima. Equivalent differentiable optimization problems and descent methods for asymmetric variational inequality problems. *Mathematical programming*, 53:99–110, 1992. (Cited on pages 105 and 114.)

Masao Fukushima. Restricted generalized nash equilibria and controlled penalty algorithm. *Computational Management Science*, 8(3):201–218, 2011. (Cited on pages 219 and 233.)

Yuan Gao and Christian Kroer. First-order methods for large-scale market equilibrium computation. *Advances in Neural Information Processing Systems*, 33:21738–21750, 2020. (Cited on page 131.)

Yuan Gao and Christian Kroer. Infinite-dimensional fisher markets: Equilibrium, duality and optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 5432–5439, 2021. (Cited on page 356.)

Michael R Garey and David S Johnson. *Computers and intractability*, volume 29. wh freeman New York, 2002. (Cited on page 2.)

Jugal Garg and Ravi Kannan. Markets with production: A polynomial time algorithm and a reduction to pure exchange. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 733–749, 2015. (Cited on page 263.)

Rahul Garg and Sanjiv Kapoor. Auction algorithms for market equilibrium. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 511–518, 2004. (Cited on page 130.)

René Gâteaux. Sur les fonctionnelles continues et les fonctionnelles analytiques. *CR Acad. Sci. Paris*, 157(325-327):65, 1913. (Cited on page 44.)

John Geanakoplos. An introduction to general equilibrium with incomplete asset markets. *Journal of mathematical economics*, 19(1-2):1–38, 1990. (Cited on pages 68, 295, 296, 297, 298, and 326.)

Ian M. Gemp, Rahul Savani, Marc Lanctot, Yoram Bachrach, Thomas W. Anthony, Richard Everett, Andrea Tacchetti, Tom Eccles, and Janos Kramar. Sample-based approximation of nash in large many-player games via gradient descent. *CoRR*, abs/2106.01285, 2021. URL https://arxiv.org/abs/2106.01285. (Cited on page 233.)

Gauthier Gidel, Hugo Berard, Gaetan Vignoud, Pascal Vincent, and Simon Lacoste-Julien. A variational inequality perspective on generative adversarial networks. *arXiv preprint arXiv:1802.10551*, 2018. (Cited on page 84.)

Benjamin J Gillen, Masayoshi Hirota, Ming Hsu, Charles R Plott, and Brian W Rogers. Divergence and convergence in scarf cycle environments: experiments and predictability in the dynamics of general equilibrium systems. *Economic Theory*, pages 1–52, 2020. (Cited on pages 66 and 129.)

Irving L Glicksberg. A further generalization of the kakutani fixed theorem, with application to nash equilibrium points. *Proceedings of the American Mathematical Society*, 3(1): 170–174, 1952. (Cited on pages 3, 39, 307, 332, and 334.)

Denizalp Goktas. An algorithmic theory of markets and their application to decentralized markets. *American Association for Artificial Intelligence Conference on Artificial Intelligence, AAAI'22*, 2022. (Cited on page iii.)

Denizalp Goktas and Amy Greenwald. Convex-concave min-max Stackelberg games. *Advances in Neural Information Processing Systems*, 34, 2021. (Cited on page iii.)

Denizalp Goktas and Amy Greenwald. Robust no-regret learning in min-max Stackelberg games, 2022a. (Cited on page iii.)

Denizalp Goktas and Amy Greenwald. Exploitability minimization in games and beyond. In *Advances in Neural Information Processing Systems*, 2022b. (Cited on pages iii, 310, and 312.)

Denizalp Goktas and Amy Greenwald. Gradient descent ascent in min-max Stackelberg games. *arXiv preprint arXiv:2208.09690*, 2022c. (Cited on page iii.)

Denizalp Goktas, Enrique Areyan Viqueira, and Amy Greenwald. A consumer-theoretic characterization of Fisher market equilibria. In Michal Feldman, Hu Fu, and Inbal Talgam-Cohen, editors, *Web and Internet Economics*, pages 334–351, Cham, 2022a. Springer International Publishing. ISBN 978-3-030-94676-0. (Cited on page iii.)

Denizalp Goktas, Enrique Areyan Viqueira, and Amy Greenwald. A consumer-theoretic characterization of fisher market equilibria. In *Web and Internet Economics: 17th International Conference, WINE 2021, Potsdam, Germany, December 14–17, 2021, Proceedings*, pages 334–351. Springer, 2022b. (Cited on pages 168, 175, 184, 185, 186, 187, 188, 199, 204, and 205.)

Denizalp Goktas, Sadie Zhao, and Amy Greenwald. Zero-sum stochastic Stackelberg games. *Advances in Neural Information Processing Systems*, 35:11658–11672, 2022c. (Cited on page iii.)

Denizalp Goktas, David C Parkes, Ian Gemp, Luke Marris, Georgios Piliouras, Romuald Elie, Guy Lever, and Andrea Tacchetti. Generative adversarial equilibrium solvers. *arXiv preprint arXiv:2302.06607*, 2023a. (Cited on pages iii, 300, and 313.)

Denizalp Goktas, Arjun Prakash, and Amy Greenwald. Convex-concave zero-sum markov stackelberg games. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023b. (Cited on page iii.)

Denizalp Goktas, Jiayi Zhao, and Amy Greenwald. Tâtonnement in homothetic Fisher markets. *arXiv preprint arXiv:2306.04890*, 2023c. (Cited on pages iii, 219, and 357.)

Denizalp Goktas, Amy Greenwald, Sadie Zhao, Alec Koppel, and Sumitra Ganesh. Efficient inverse multiagent learning. In *The Twelfth International Conference on Learning Representations*, 2024. (Cited on page iii.)

Seth Copen Goldstein, **Goktas, Denizalp**, Miles Conn, Shanmukha Phani Teja Pitchuka, Mohammed Sameer, Maya Shah, Hefei Tu Colin Swett, Shrinath Viswanathan, and Jessica Xiao. Bolt: Building on local trust to solve lending market failure. *University Blockchain Research Initiative 2020, UBRI'20*, 2020. (Cited on page iii.)

Noah Golowich, Sarath Pattathil, and Constantinos Daskalakis. Tight last-iterate convergence rates for no-regret learning in multi-player games. *Advances in neural information processing systems*, 33:20766–20778, 2020a. (Cited on pages 84, 230, and 231.)

Noah Golowich, Sarath Pattathil, Constantinos Daskalakis, and Asuman E. Ozdaglar. Last iterate is slower than averaged iterate in smooth convex-concave saddle point problems. *CoRR*, abs/2002.00057, 2020b. URL https://arxiv.org/abs/2002.00057. (Cited on pages 87 and 88.)

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf. (Cited on pages 75 and 329.)

Eduard Gorbunov, Nicolas Loizou, and Gauthier Gidel. Extragradient method: O (1/k) last-iterate convergence for monotone variational inequalities and connections with cocoercivity. In *International Conference on Artificial Intelligence and Statistics*, pages 366–402. PMLR, 2022a. (Cited on page 84.)

Eduard Gorbunov, Adrien Taylor, and Gauthier Gidel. Last-iterate convergence of optimistic gradient method for monotone variational inequalities. *Advances in neural information processing systems*, 35:21858–21870, 2022b. (Cited on page 84.)

Lawrence M Graves. Riemann integration and taylor's theorem in general analysis. *Transactions of the American Mathematical Society*, 29(1):163–177, 1927. (Cited on pages 177 and 209.)

Giuseppe Grioli. Proprietà variazionali nella meccanica dei continui. nota i. *Atti della Accademia Nazionale dei Lincei. Classe di Scienze Fisiche, Matematiche e Naturali. Rendiconti*, 54(5):755–760, 1973. (Cited on page 75.)

Frank H Hahn. Gross substitutes and the dynamic stability of general equilibrium. *Econometrica (pre-1986)*, 26(1):169, 1958. (Cited on page 219.)

Zhu Han, Dusit Niyato, Walid Saad, Tamer Basar, and Are Hjorungnes. *Game theory in wireless and communication networks: Theory, models, and applications*, volume 9780521196963. Cambridge University Press, United Kingdom, January 2011. ISBN 9780521196963. doi: 10.1017/CBO9780511895043. Publisher Copyright: Cambridge University Press 2012. (Cited on page 218.)

Charles R. Harris, K. Jarrod Millman, Stéfan J van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernandez del Rio, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585:357–362, 2020. doi: 10.1038/s41586-020-2649-2. (Cited on pages 191 and 352.)

Oliver D Hart. On the optimality of equilibrium when the market structure is incomplete. *Journal of economic theory*, 11(3):418–443, 1975. (Cited on page 298.)

Sergiu Hart and Andreu Mas-Colell. Uncoupled dynamics do not lead to nash equilibrium. *The American Economic Review*, 93(5):1830–1836, 2003. ISSN 00028282. URL `http://www.jstor.org/stable/3132156`. (Cited on pages 230 and 240.)

Philip Hartman and Guido Stampacchia. On some non-linear elliptic differential-functional equations. 1966. (Cited on page 75.)

Bingsheng He, Shengjie Xu, and Xiaoming Yuan. On convergence of the arrow–hurwicz method for saddle point problems. *Journal of Mathematical Imaging and Vision*, 64(6): 662–671, 2022. (Cited on pages 81, 83, and 84.)

Yiran He. Solvability of the minty variational inequality. *Journal of Optimization Theory and Applications*, 174:686–692, 2017. (Cited on pages 81, 145, and 148.)

Tom Hennigan, Trevor Cai, Tamara Norman, and Igor Babuschkin. Haiku: Sonnet for JAX, 2020. URL `http://github.com/deepmind/dm-haiku`. (Cited on page 352.)

P Jean-Jacques Herings. A globally and universally stable price adjustment process. *Journal of Mathematical Economics*, 27(2):163–193, 1997. (Cited on page 67.)

A. Heusinger and C. Kanzow. Relaxation methods for generalized nash equilibrium problems with inexact line search. *Journal of Optimization Theory and Applications*, 143 (1):159–183, 2009. URL `https://EconPapers.repec.org/RePEc:spr:joptap:v:143:y:2009:i:1:d:10.1007_s10957-009-9553-0`. (Cited on page 232.)

Benjamin F. Hobbs and J. S. Pang. Nash-cournot equilibria in electric power markets with piecewise linear demand functions and joint constraints. *Operations Research*, 55 (1):113–127, 2007. doi: 10.1287/opre.1060.0342. URL `https://doi.org/10.1287/opre.1060.0342`. (Cited on page 218.)

Kevin Huang and Shuzhong Zhang. An approximation-based regularized extra-gradient method for monotone variational inequalities. *arXiv preprint arXiv:2210.04440*, 2022. (Cited on page 83.)

Kevin Huang and Shuzhong Zhang. Beyond monotone variational inequalities: Solution methods and iteration complexities. *arXiv preprint arXiv:2304.04153*, 2023. (Cited on pages 8, 71, 88, 105, 238, 239, and 255.)

Kevin XD Huang and Jan Werner. Asset price bubbles in arrow-debreu and sequential equilibrium. *Economic Theory*, 15:253–278, 2000. (Cited on page 299.)

Kevin XD Huang and Jan Werner. Implementing arrow-debreu equilibria by trading infinitely-lived securities. *Economic Theory*, 24:603–622, 2004. (Cited on page 321.)

J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science and Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55. (Cited on pages 191 and 352.)

Leonid Hurwicz. Optimality and informational efficiency in resource allocation processes. *Mathematical methods in the social sciences*, 1960. (Cited on page 23.)

Leonid Hurwicz. On informationally decentralized systems. *Decision and organization: A volume in Honor of J. Marschak*, 1972. (Cited on page 23.)

Leonid Hurwicz. Outcome functions yielding walrasian and lindahl allocations at nash equilibrium points. *The Review of Economic Studies*, 46(2):217–225, 1979. (Cited on page 23.)

UC IEOR. The linear complementarity problem, lemke algorithm, perturbation, and the complexity class ppad. 2011. (Cited on page 127.)

Fumitada Itakura and Shuzo Saito. Analysis synthesis telephony based on the maximum likelihood method. In *Proceedings of the 6th International Congress on Acoustics*, pages C17–C20, 1968. (Cited on page 358.)

Alexey F Izmailov and Mikhail V Solodov. On error bounds and newton-type methods for generalized nash equilibrium problems. *Computational Optimization and Applications*, 59 (1):201–218, 2014. (Cited on pages 219 and 232.)

Kamal Jain, Vijay V Vazirani, and Yinyu Ye. Market equilibria for homothetic, quasi-concave utilities and economies of scale in production. In *SODA*, volume 5, pages 63–71, 2005. (Cited on pages 65, 129, 168, and 219.)

Shweta Jain, Narayanaswamy Balakrishnan, Yadati Narahari, Saiful A Hussain, and Nyuk Yoong Voo. Constrained tâtonnement for fast and incentive compatible distributed demand management in smart grids. In *Proceedings of the fourth international conference on Future energy systems*, pages 125–136, 2013. (Cited on pages 68, 128, and 130.)

Xuegang (Jeff) Ban, Maged Dessouky, Jong-Shi Pang, and Rong Fan. A general equilibrium model for transportation systems with e-hailing services and flow congestion. *Transportation Research Part B: Methodological*, 129:273–304, 2019. ISSN 0191-2615. doi: https://doi.org/10.1016/j.trb.2019.08.012. URL https://www.sciencedirect.com/science/article/pii/S0191261518309044. (Cited on page 218.)

Geoffrey Alexander Jehle. *Advanced microeconomic theory*. Pearson Education India, 2001. (Cited on page 166.)

Wei Jing-Yuan and Yves Smeers. Spatial oligopolistic electricity models with cournot generators and regulated transmission prices. *Operations Research*, 47(1):102–112, 1999. doi: 10.1287/opre.47.1.102. URL https://pubsonline.informs.org/doi/abs/10.1287/opre.47.1.102. (Cited on page 218.)

Michael I. Jordan, Tianyi Lin, and Manolis Zampetakis. First-order algorithms for nonlinear generalized nash equilibrium problems, 2022. URL https://arxiv.org/abs/2204.03132. (Cited on pages 219 and 233.)

James M Joyce. Kullback-leibler divergence. In *International encyclopedia of statistical science*, pages 720–722. Springer, 2011. (Cited on page 163.)

Kenneth L Judd. Projection methods for solving aggregate growth models. *Journal of Economic Theory*, 58(2):410–452, December 1992. ISSN 0022-0531. doi: 10.1016/0022-0531(92)90061-L. URL https://www.sciencedirect.com/science/article/pii/002205319290061L. (Cited on page 346.)

Sham M Kakade, Michael Kearns, and Luis E Ortiz. Graphical economics. In *COLT*, pages 17–32. Springer, 2004. (Cited on page 131.)

Shizuo Kakutani. A generalization of brouwer's fixed point theorem. *Duke Mathematical Journal*, 7:457–459, 1941. (Cited on pages 3 and 39.)

Kazuya Kamiya. A globally stable price adjustment process. *Econometrica: Journal of the Econometric Society*, pages 1481–1485, 1990. (Cited on page 67.)

Christian Kanzow. On the multiplier-penalty-approach for quasi-variational inequalities. *Mathematical Programming*, 160(1):33–63, 2016. (Cited on pages 219 and 233.)

Christian Kanzow and Daniel Steck. Augmented lagrangian methods for the solution of generalized nash equilibrium problems. *SIAM Journal on Optimization*, 26(4):2034–2058, 2016. (Cited on page 219.)

Christian Kanzow and Daniel Steck. Augmented lagrangian and exact penalty methods for quasi-variational inequalities. *Computational Optimization and Applications*, 69(3):801–824, 2018. (Cited on pages 219 and 233.)

Bruce M Kapron and Koosha Samieefar. The computational complexity of variational inequalities and applications in game theory. *arXiv preprint arXiv:2411.04392*, 2024. (Cited on pages 96 and 111.)

Jinill Kim and Sunghyun Henry Kim. Spurious welfare reversals in international business cycle models. *journal of International Economics*, 60(2):471–500, 2003. (Cited on pages 2 and 5.)

Galina M Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976. (Cited on pages 71, 83, and 241.)

Jacek B. Krawczyk. Coupled constraint nash equilibria in environmental games. *Resource and Energy Economics*, 27(2):157–181, 2005. ISSN 0928-7655. doi: https://doi.org/10.1016/j.reseneeco.2004.08.001. URL `https://www.sciencedirect.com/science/article/pii/S0928765504000661`. (Cited on page 218.)

Jacek B. Krawczyk and Stanislav Uryasev. Relaxation algorithms to find nash equilibria with economic applications. *Environmental Modeling and Assessment*, 5(1):63–73, 2000. doi: 10.1023/A:1019097208499. This revised version was published online in July 2006 with corrections to the Cover Date. (Cited on page 232.)

HW Kuhn and AW Tucker. Proceedings of 2nd berkeley symposium, 1951. (Cited on pages 57 and 228.)

Jerzy Kyparisis. On uniqueness of kuhn-tucker multipliers in nonlinear programming. *Mathematical Programming*, 32:242–246, 1985. (Cited on page 58.)

Jerzy Kyparisis and Anthony V Fiacco. Generalized convexity and concavity of the optimal value function in nonlinear programming. *Mathematical Programming*, 39:285–304, 1987. (Cited on page 60.)

Peter D Lax. The flowering of applied mathematics in america. *Siam Review*, 31(4):533–541, 1989. (Cited on page 2.)

Stefanos Leonardos, Barnabé Monnot, Daniël Reijsbergen, Efstratios Skoulakis, and Georgios Piliouras. Dynamical analysis of the eip-1559 ethereum fee market. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, pages 114–126, 2021. (Cited on pages 68, 128, and 130.)

Jonathan Levin. Lecture notes on consumer theory, October 2004. (Cited on pages 172 and 203.)

Tianyi Lin and Michael I. Jordan. Perseus: A simple high-order regularization method for variational inequalities. *ArXiv*, abs/2205.03202, 2022. URL `https://api.semanticscholar.org/CorpusID:248562799`. (Cited on page 84.)

Tianyi Lin, Chi Jin, and Michael Jordan. On gradient descent ascent for nonconvex-concave minimax problems. In *International Conference on Machine Learning*, pages 6083–6093. PMLR, 2020. (Cited on pages 50, 51, 300, 316, and 342.)

John Lintner. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. In *Stochastic optimization models in finance*, pages 131–155. Elsevier, 1975. (Cited on page 297.)

Jacques-Louis Lions and Guido Stampacchia. Variational inequalities. *Communications on pure and applied mathematics*, 20(3):493–519, 1967. (Cited on pages 3, 75, and 76.)

Mingrui Liu, Hassan Rafique, Qihang Lin, and Tianbao Yang. First-order convergence theory for weakly-convex-weakly-concave min-max problems. *The Journal of Machine Learning Research*, 22(1):7651–7684, 2021. (Cited on pages 84, 86, and 310.)

Yulin Liu, Yuxuan Lu, Kartik Nayak, Fan Zhang, Luyao Zhang, and Yinhong Zhao. Empirical analysis of eip-1559: Transaction fees, waiting time, and consensus security. *arXiv preprint arXiv:2201.05574*, 2022. (Cited on pages 68, 128, and 130.)

Edward Lockhart, Marc Lanctot, Julien Perolat, Jean-Baptiste Lespiau, Dustin Morrill, Finbarr Timbers, and Karl Tuyls. Computing approximate equilibria in sequential adversarial games by exploitability descent. *CoRR*, abs/1903.05614, 2019. URL http://arxiv.org/abs/1903.05614. (Cited on page 233.)

Nicolas Loizou, Hugo Berard, Gauthier Gidel, Ioannis Mitliagkas, and Simon Lacoste-Julien. Stochastic gradient descent-ascent and consensus optimization for smooth games: Convergence analysis under expected co-coercivity. *CoRR*, abs/2107.00052, 2021. URL https://arxiv.org/abs/2107.00052. (Cited on page 84.)

Haihao Lu. "relative continuity" for non-lipschitz nonsmooth convex optimization using stochastic (or deterministic) mirror descent. *INFORMS Journal on Optimization*, 1(4): 288–303, 2019. (Cited on pages 69 and 138.)

Robert E Lucas Jr and Edward C Prescott. Investment under uncertainty. *Econometrica: Journal of the Econometric Society*, pages 659–681, 1971. (Cited on page 323.)

Michael Magill and Martine Quinzii. Infinite horizon incomplete markets. *Econometrica: Journal of the Econometric Society*, pages 853–880, 1994. (Cited on pages 12, 299, 300, and 321.)

Michael Magill and Martine Quinzii. *Theory of incomplete markets*, volume 1. Mit press, 2002. (Cited on pages 297 and 298.)

Michael Magill and Wayne Shafer. Incomplete markets. *Handbook of mathematical economics*, 4:1523–1614, 1991. (Cited on page 297.)

Pouria Mahdavinia, Yuyang Deng, Haochuan Li, and Mehrdad Mahdavi. Tight analysis of extra-gradient and optimistic gradient methods for nonconvex minimax problems. *Advances in Neural Information Processing Systems*, 35:31213–31225, 2022. (Cited on pages 244 and 246.)

Jacob Makar-Limanov, Arjun Prakash, Denizalp Goktas, Nora Ayanian, and Amy Greenwald. Sta-rlhf: Stackelberg aligned reinforcement learning with human feedback. In *Coordination and Cooperation for Multi-Agent Reinforcement Learning Methods Workshop*, 2024. (Cited on page iii.)

Lilia Maliar, Serguei Maliar, and Pablo Winant. Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, 122:76–101, 2021. (Cited on page 347.)

Alfred Marshall. *Principles of economics: An introductory volume*. Macmillan, 1910. (Cited on page 15.)

Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Number 9780195102680 in OUP Catalogue. Oxford University Press, 1995. ISBN ARRAY(0x4cf9c5c0). URL `https://ideas.repec.org/b/oxp/obooks/9780195102680.html`. (Cited on pages 166, 179, 203, 205, and 298.)

Eric Maskin and Jean Tirole. Markov perfect equilibrium: I. observable actions. *Journal of Economic Theory*, 100(2):191–219, 2001. (Cited on page 304.)

Lionel W McKenzie. On the existence of general equilibrium for a competitive market. *Econometrica: journal of the Econometric Society*, pages 54–71, 1959. (Cited on page 219.)

R Mehra and EC Prescott. Recursive competitive equilibria and capital asset pricing. *Essays in Financial Economics*, 1977. (Cited on pages 300 and 326.)

Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70 (2):583–601, 2002. (Cited on page 178.)

John Stuart Mill. Principles of political economy. *U of Toronto P*, 1965. (Cited on page 15.)

George J Minty. On the generalization of a direct method of the calculus of variations. 1967. (Cited on pages 69 and 78.)

Benjamin H Mitra-Kahn. General equilibrium theory, its history and its relation (if any) to the market economy. *City University, London The New School for Social Research, New York*, 2005. (Cited on pages 17, 21, and 22.)

Franco Modigliani and Merton H Miller. The cost of capital, corporation finance and the theory of investment. *The American economic review*, 48(3):261–297, 1958. (Cited on page 297.)

Jan Mossin. Equilibrium in a capital asset market. *Econometrica: Journal of the econometric society*, pages 768–783, 1966. (Cited on page 297.)

Katta G Murty and Santosh N Kabadi. Some np-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39:117–129, 1987. (Cited on page 328.)

Koichi Nabetani, Paul Tseng, and Masao Fukushima. Parametrized variational inequality approaches to generalized nash equilibrium problems with shared constraints. *Comput. Optim. Appl.*, 48(3):423–452, 4 2011. ISSN 0926-6003. doi: 10.1007/s10589-009-9256-3. URL https://doi.org/10.1007/s10589-009-9256-3. (Cited on page 232.)

Anna Nagurney. *Network economics*. Wiley Online Library, 2009. (Cited on page 117.)

John F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950a. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1907266. (Cited on page 309.)

John F. Nash. Equilibrium points in <i>n</i>-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950b. doi: 10.1073/pnas.36.1.48. URL https://www.pnas.org/doi/abs/10.1073/pnas.36.1.48. (Cited on pages 219, 223, and 306.)

Angelia Nedić and Asuman Ozdaglar. Approximate primal solutions and rate analysis for dual subgradient methods. *SIAM Journal on Optimization*, 19(4):1757–1780, 2009. (Cited on page 289.)

Angelia Nedic and Asuman Ozdaglar. Subgradient methods for saddle-point problems. *Journal of optimization theory and applications*, 142(1):205–228, 2009. (Cited on pages 229, 241, and 289.)

Arkadi Nemirovski. Prox-method with rate of convergence o (1/t) for variational inequalities with lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004. (Cited on pages 84, 87, 88, 241, 242, and 243.)

Arkadij Semenovic Nemirovskij and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983. (Cited on page 85.)

Yurii Nesterov. Introductory lectures on convex programming volume i: Basic course. *Lecture notes*, 3(4):5, 1998. (Cited on pages 76 and 83.)

Hukukane Nikaido and Kazuo Isoda. Note on non-cooperative convex games. *Pacific Journal of Mathematics*, 5(S1):807–815, 1955. (Cited on page 240.)

Hukukane Nikaidô and Hirofumi Uzawa. Stability and non-negativity in a walrasian tatonnement process. *International Economic Review*, 1(1):50–59, 1960. (Cited on page 66.)

Noam Nisan and Tim Roughgarden. *Algorithmic Game Theory*. Cambridge University Press, 2007. doi: 10.1017/CBO9780511800481. (Cited on pages 128, 129, and 166.)

Evgeni Alekseevich Nurminskii. The quasigradient method for the solving of the nonlinear programming problems. *Cybernetics*, 9(1):145–150, 1973. (Cited on page 50.)

Austine Efut Ofem, Akindele Adebayo Mebawondu, Godwin Chidi Ugwunnadi, Huseyin Isik, and Ojen Kumar Narain. A modified subgradient extragradient algorithm-type for solving quasimonotone variational inequality problems with applications. *Journal of Inequalities and Applications*, 2023(1):73, 2023. (Cited on page 84.)

Dario Paccagnan, Basilio Gentile, Francesca Parise, Maryam Kamgarpour, and John Lygeros. Distributed computation of generalized nash equilibria in quadratic aggregative games with affine coupling constraints. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 6123–6128. IEEE, 2016. (Cited on page 219.)

Jong-Shi Pang and Masao Fukushima. Quasi-variational inequalities, generalized nash equilibria, and multi-leader-follower games. *Computational Management Science*, 2(1): 21–56, 2005. (Cited on pages 219 and 233.)

Jong-Shi Pang, Gesualdo Scutari, Francisco Facchinei, and Chaoxiong Wang. Distributed power allocation with rate constraints in gaussian parallel interference channels. *IEEE Transactions on Information Theory*, 54(8):3471–3489, 2008. (Cited on page 218.)

Christos H Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and system Sciences*, 48(3):498–532, 1994. (Cited on pages 2 and 3.)

Christos H Papadimitriou and Mihalis Yannakakis. An impossibility theorem for price-adjustment mechanisms. *Proceedings of the National Academy of Sciences*, 107(5):1854–1859, 2010. (Cited on pages 68, 69, 128, 129, 132, 143, and 357.)

I. F. Pearce and J. Wise. On the uniqueness of competitive equilibrium: Part i, unbounded demand. *Econometrica*, 41(5):817–828, 1973. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1913808. (Cited on page 219.)

Leonid Denisovich Popov. A modification of the arrow-hurwitz method of search for saddle points. *Mat. Zametki*, 28(5):777–784, 1980. (Cited on pages 83 and 87.)

Edward C Prescott and Robert E Lucas. A note on price systems in infinite dimensional space. *International Economic Review*, pages 416–422, 1972. (Cited on pages 3 and 356.)

Edward C. Prescott and Rajnish Mehra. Recursive competitive equilibrium: The case of homogeneous households. *Econometrica*, 48(6):1365–1379, 1980. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1912812. (Cited on pages 300 and 326.)

JD Pryce. R. tyrell rockafellar, convex analysis (princeton university press, 1970), xviii+ 451 pp. *Proceedings of the Edinburgh Mathematical Society*, 18(4):339–339, 1973. (Cited on page 45.)

Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014. (Cited on page 311.)

Roy Radner. Competitive equilibrium under uncertainty. *Econometrica: Journal of the Econometric Society*, pages 31–58, 1968. (Cited on pages 19, 20, and 297.)

Roy Radner. Existence of equilibrium of plans, prices, and price expectations in a sequence of markets. *Econometrica: Journal of the Econometric Society*, pages 289–303, 1972. (Cited on pages 18, 19, 297, 320, and 326.)

Roy Radner. Rational expectations equilibrium: Generic existence and the information revealed by prices. *Econometrica: Journal of the Econometric Society*, pages 655–678, 1979. (Cited on page 298.)

Rockafellar Ralph Tyrell. *Convex Analysis : (PMS-28).* Princeton Paperbacks. Princeton University Press, 1997. ISBN 9780691015866. URL https://search.ebscohost.com/login.aspx?direct=true&db=e000xna&AN=969057&site=ehost-live&scope=site&authtype=ip,sso&custid=rock. (Cited on page 49.)

Jacob Ramskov and Jesper Munksgaard. Elasticities–a theoretical introduction. *Balmorel Project*, 2001. (Cited on page 175.)

John Randolph, Denizalp Goktas, and Amy Greenwald. Banzhaf power in hierarchical games. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'24)*, 2024. (Cited on page iii.)

Daniël Reijsbergen, Shyam Sridhar, Barnabé Monnot, Stefanos Leonardos, Stratis Skoulakis, and Georgios Piliouras. Transaction fees on a honeymoon: Ethereum's eip-1559 one month later. In *2021 IEEE International Conference on Blockchain (Blockchain)*, pages 196–204. IEEE, 2021. (Cited on pages 68, 128, and 130.)

David Ricardo. *The first six chapters of the principles of political economy and taxation of David Ricardo, 1817*. Macmillan and Company, 1895. (Cited on page 15.)

Jeremy Rifkin. *The third industrial revolution: how lateral power is transforming energy, the economy, and the world*. Macmillan, 2011. (Cited on page 1.)

R Tyrrell Rockafellar. Saddle-points and convex analysis. *Differential games and related topics*, 109, 1971. (Cited on page 58.)

R Tyrrell Rockafellar and Roger J-B Wets. *Variational analysis*, volume 317. Springer Science and Business Media, 2009. (Cited on pages 49 and 60.)

J. B. Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, 33(3):520–534, 1965. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/1911749. (Cited on pages 226 and 232.)

Paul A Samuelson. The pure theory of public expenditure. *The review of economics and statistics*, 36(4):387–389, 1954. (Cited on page 21.)

Thomas J Sargent and Lars Ljungqvist. Recursive macroeconomic theory. *Massachusetss Institute of Technology*, 2000. (Cited on pages 68, 249, 298, and 299.)

Maxime Sauzet. Projection Methods via Neural Networks for Continuous-Time Models, December 2021. URL https://papers.ssrn.com/abstract=3981838. (Cited on page 347.)

Herbert Scarf. Some examples of global instability of the competitive equilibrium. *International Economic Review*, 1(3):157–172, 1960. ISSN 00206598, 14682354. URL http://www.jstor.org/stable/2556215. (Cited on pages 67, 155, 353, and 357.)

Herbert Scarf. The approximation of fixed points of a continuous mapping. *SIAM Journal on Applied Mathematics*, 15(5):1328–1343, 1967a. (Cited on pages 1 and 66.)

Herbert E Scarf. On the computation of equilibrium prices. 1967b. (Cited on page 66.)

Herbert E Scarf. The computation of equilibrium prices: an exposition. *Handbook of mathematical economics*, 2:1007–1061, 1982. (Cited on page 66.)

Herbert E Scarf and Terje Boye Hansen. The computation of economic equilibria. *A Cowles Foundation Monograph*, 1973. (Cited on page 66.)

Dane A. Schiro, Jong-Shi Pang, and Uday V. Shanbhag. On the solution of affine generalized nash equilibrium problems with shared constraints by lemke's method. *Mathematical Programming*, 142(1-2):1–46, 2013. doi: 10.1007/s10107-012-0558-3. This work was based on research partially supported by the National Science Foundation under grant CMMI-0969600 and the Department of Energy under grant DOE DE-SC0003879. (Cited on page 232.)

Hermann Amandus Schwarz. Über einige abhängigkeiten zwischen den werteigen einer function und deren ableitungen. *Journal für die reine und angewandte Mathematik*, 98: 159–181, 1884. In German. Expands the inequality in the context of integrals and inner product spaces. (Cited on page 90.)

Lloyd Shapley and Martin Shubik. Trade using one commodity as a means of payment. *Journal of political economy*, 85(5):937–968, 1977. (Cited on page 275.)

Lloyd S Shapley. Stochastic games. *Proceedings of the national academy of sciences*, 39(10): 1095–1100, 1953. (Cited on page 302.)

Ronald William Shephard. *Theory of cost and production functions*. Princeton University Press, 2015. (Cited on pages 172, 178, 197, 199, and 210.)

David Sheppard, Myles McCormick, Anjli Raval, Derek Brower, and Hudson Lockett. Us oil price below zero for first time in history. *Financial Times*, 21(4):1–8, 2020. (Cited on page 321.)

David Silver, Guy Lever, Nicolas Heess, Thomas Degris, Daan Wierstra, and Martin Riedmiller. Deterministic policy gradient algorithms. In *International conference on machine learning*, pages 387–395. Pmlr, 2014. (Cited on page 339.)

Steve Smale. A convergent process of price adjustment and global newton methods. *Journal of Mathematical Economics*, 3(2):107–120, 1976. (Cited on page 67.)

Adam Smith. *The wealth of nations [1776]*, volume 11937. na, 1937. (Cited on page 15.)

Mikhail V Solodov and Benar F Svaiter. A hybrid approximate extragradient–proximal point algorithm using the enlargement of a maximal monotone operator. *Set-Valued Analysis*, 7(4):323–345, 1999. (Cited on page 83.)

Hugo Sonnenschein. Market excess demand functions. *Econometrica: Journal of the Econometric Society*, pages 549–563, 1972. (Cited on page 123.)

Oliver Stein and Nathan Sudermann-Merx. The noncooperative transportation problem and linear generalized Nash games. *European Journal of Operational Research*, 266(2): 543–553, 2018. doi: 10.1016/j.ejor.2017.10.00. URL `https://ideas.repec.org/a/eee/ejores/v266y2018i2p543-553.html`. (Cited on page 218.)

Hyung Ju Suh, Max Simchowitz, Kaiqing Zhang, and Russ Tedrake. Do differentiable simulators give better policy gradients? In *International Conference on Machine Learning*, pages 20668–20696. PMLR, 2022. (Cited on page 317.)

Masayuki Takahashi. Equilibrium points of stochastic non-cooperative n-person games. *Journal of Science of the Hiroshima University, Series AI (Mathematics)*, 28(1):95–99, 1964. (Cited on pages 11, 299, and 302.)

Yoshihiro Tanaka. Nonsmooth optimization for production theory. *Hokkaido University, Sapporo*, 2008. (Cited on pages 172, 197, and 199.)

Tatiana Tatarenko and Maryam Kamgarpour. Learning generalized nash equilibria in a class of convex games. *IEEE Transactions on Automatic Control*, 64(4):1426–1439, 2018. (Cited on page 219.)

John B Taylor and Michael Woodford. *Handbook of macroeconomics*, volume 1. Elsevier, 1999. (Cited on pages 68 and 299.)

Duong Viet Thong, Dang Van Hieu, and Themistocles M Rassias. Self adaptive inertial subgradient extragradient algorithms for solving pseudomonotone variational inequality problems. *Optimization Letters*, 14(1):115–144, 2020. (Cited on page 84.)

S. Uryas'ev and R.Y. Rubinstein. On relaxation algorithms in computation of noncooperative equilibria. *IEEE Transactions on Automatic Control*, 39(6):1263–1267, 1994. doi: 10.1109/9.293193. (Cited on page 232.)

Hirofumi Uzawa. Walras' tatonnement in the theory of exchange. *The Review of Economic Studies*, 27(3):182–194, 1960. (Cited on pages 15, 65, and 353.)

Victor Amarachi Uzor, Timilehin Opeyemi Alakoya, Oluwatosin Temitope Mewomo, and Aviv Gibali. Solving quasimonotone and non-monotone variational inequalities. *Mathematical Methods of Operations Research*, 98(3):461–498, 2023. (Cited on page 84.)

Gerard van der Laan, AJJ Talman, et al. A convergent price adjustment process. *Economics Letters*, 23(2):119–123, 1987. (Cited on page 67.)

Guido Van Rossum and Fred L Drake Jr. *Python tutorial*. Centrum voor Wiskunde en Informatica Amsterdam, The Netherlands, 1995. (Cited on pages 191 and 352.)

Stephen A Vavasis. Complexity issues in global optimization: a survey. *Handbook of global optimization*, pages 27–41, 1995. (Cited on pages 230 and 248.)

Jean-Philippe Vial. Strong and weak convexity of sets and functions. *Mathematics of Operations Research*, 8(2):231–259, 1983. (Cited on page 50.)

Xavier Vives. Small income effects: A marshallian theory of consumer surplus and downward sloping demand. *The Review of Economic Studies*, 54(1):87–103, 1987. ISSN 00346527, 1467937X. URL http://www.jstor.org/stable/2297448. (Cited on page 172.)

Anna Von Heusinger and Christian Kanzow. Optimization reformulations of the generalized nash equilibrium problem using nikaido-isoda-type functions. *Computational Optimization and Applications*, 43(3):353–377, 2009. (Cited on pages 219, 245, 246, and 247.)

Anna Von Heusinger, Christian Kanzow, and Masao Fukushima. Newton's method for computing a normalized equilibrium in the generalized nash game through fixed point formulation. *Mathematical Programming*, 132(1):99–123, 2012. (Cited on pages 219 and 232.)

John von Neumann. Zur theorie der gesellschaftsspiele. *Mathematische annalen*, 100(1): 295–320, 1928. (Cited on page 241.)

Abraham Wald. Statistical decision functions which minimize the maximum risk. *Annals of Mathematics*, pages 265–280, 1945. (Cited on page 312.)

Leon Walras. *Elements de l'economie politique pure, ou, Theorie de la richesse sociale*. F. Rouge, 1896. (Cited on pages 15, 64, 65, 115, 116, 132, 133, 217, 294, and 353.)

WY Wang and BB Ma. Extragradient type projection algorithm for solving quasimonotone variational inequalities. *Optimization*, 73(8):2639–2655, 2024. (Cited on page 84.)

Graham R Wood and BP Zhang. Estimation of the lipschitz constant of a function. *Journal of Global Optimization*, 8:91–103, 1996. (Cited on page 160.)

Fang Wu and Li Zhang. Proportional response dynamics leads to market equilibrium. In *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing*, STOC '07, page 354–363, New York, NY, USA, 2007. Association for Computing Machinery.

ISBN 9781595936318. doi: 10.1145/1250790.1250844. URL https://doi.org/10.1145/1250790.1250844. (Cited on page 131.)

Mihalis Yannakakis. Equilibria, fixed points, and complexity classes. *Computer Science Review*, 3(2):71–85, 2009. (Cited on page 3.)

Peng Yi and Lacra Pavel. A distributed primal-dual algorithm for computation of generalized nash equilibria via operator splitting methods. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 3841–3846. IEEE, 2017. (Cited on page 219.)

I Zang and M Avriel. On functions whose local minima are global. *Journal of Optimization Theory and Applications*, 16:183–190, 1975. (Cited on page 80.)

Hui Zhang and Yu-Hong Dai. Mirror frameworks for relatively lipschitz and monotone-like variational inequalities. *Journal of the Operations Research Society of China*, pages 1–31, 2023. (Cited on pages 8 and 71.)

Li Zhang. Proportional response dynamics in the Fisher market. *Theor. Comput. Sci.*, 412(24):2691–2698, May 2011. ISSN 0304-3975. doi: 10.1016/j.tcs.2010.06.021. URL https://doi.org/10.1016/j.tcs.2010.06.021. (Cited on page 131.)

Siqi Zhang and Niao He. On the convergence rate of stochastic mirror descent for nonsmooth nonconvex optimization. *arXiv preprint arXiv:1806.04781*, 2018. (Cited on page 85.)

Jiayi Zhao, Denizalp Goktas, and Amy Greenwald. Fisher markets with social influence. *Proceedings of the AAAI Conference on Artificial Intelligence*, 37(5):5900–5909, Jun. 2023. doi: 10.1609/aaai.v37i5.25731. URL https://ojs.aaai.org/index.php/AAAI/article/view/25731. (Cited on page iii.)

Zhengyuan Zhou, Panayotis Mertikopoulos, Aris L Moustakas, Nicholas Bambos, and Peter Glynn. Mirror descent learning in continuous games. In *2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, pages 5776–5783. IEEE, 2017. (Cited on page 236.)