

**Dynamic Strategic Interactions:
Analysis and Mechanism Design**

by

Utku Ozan Candogan

B.S., Electrical and Electronics Engineering, Bilkent University (2007)
M.S., Electrical Engineering and Computer Science, Massachusetts Institute of
Technology (2009)


Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2013



© Massachusetts Institute of Technology 2013. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
May 22, 2013

Certified by
✓ A ✓
Asuman Ozdaglar
Professor
Thesis Supervisor

Certified by
✓
Pablo A. Parrilo
Professor
Thesis Supervisor

Accepted by
/ Leslie A. Kolodziejski
Professor
Chair, Department Committee on Graduate Students

Dynamic Strategic Interactions: Analysis and Mechanism Design

by

Utku Ozan Candogan

Submitted to the Department of Electrical Engineering and Computer Science
on May 22, 2013, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Modern systems, such as engineering systems with autonomous entities, markets, and financial networks, consist of self-interested agents with potentially conflicting objectives. These agents interact in a dynamic manner, modifying their strategies over time to improve their payoffs. The presence of self-interested agents in such systems, necessitates *the analysis* of the impact of multi-agent decision making on the overall system, and *the design* of new systems with improved performance guarantees.

Motivated by this observation, in the first part of this thesis we focus on fundamental structural properties of games, and exploit them to provide a new framework for analyzing the limiting behavior of strategy update rules in various game-theoretic settings. In the second part, we investigate the design problem of an auctioneer who uses iterative multi-item auctions for efficient allocation of resources.

More specifically, in the first part of the thesis we focus on potential games, a special class of games with desirable equilibrium and dynamic properties, and analyze their preference structure. Exploiting this structure we obtain a decomposition of arbitrary games into three components, which we refer to as the potential, harmonic, and nonstrategic components. Intuitively, the potential component of a game captures interactions that can equivalently be represented as a common interest game, while the harmonic part represents conflicts between the interests of the players. We make this intuition precise by studying the properties of these two components, and establish that indeed they have quite distinct and remarkable characteristics. The decomposition also allows us to approximate a given game with a potential game. We show that the set of approximate equilibria of an arbitrary game can be characterized through the equilibria of a potential game that approximates it.

The decomposition provides a valuable tool for the analysis of dynamics in games. Earlier literature established that many natural strategy update rules converge to a Nash equilibrium in potential games. We show that games that are close to a potential game exhibit similar properties. In particular, we focus on three commonly studied discrete-time update rules (better/best response, logit response, and discrete-time fictitious play dynamics), and establish that in near-potential games, the limiting behavior of these update rules can be characterized by an approximate equilibrium set, size of which is proportional to the distance of the original game from a potential game. Since a close potential game to a given game can be systematically found via decomposition, our results suggest a systematic framework for studying the limiting behavior of adaptive dynamics in arbitrary finite strategic form games: the limiting behavior of dynamics in a given game can be characterized by first approximating this game with a potential game, and then analyzing the limiting behavior of dynamics in the potential game.

In the second part of the thesis, we change our focus to implementing efficient outcomes in multi-agent settings by using simple mechanisms. In particular, we develop novel efficient iterative auction formats for multi-item environments, where items exhibit value complementarities/substitutabilities. We obtain our results by focusing on a special class of value functions, which we refer to as graphical valuations. These valuations are not fully general, but importantly they capture value complementarity/substitutability in important practical settings, while allowing for a compact representation of the value functions.

We start our analysis by first analyzing how the special structure of graphical valuations can be exploited to design simple iterative auction formats. We show that in settings where the underlying value graph is a tree (and satisfies an additional technical condition), a Walrasian equilibrium always exists (even in the presence of value complementarities). Using this result we provide a linear programming formulation of the efficient allocation problem for this class of valuations. Additionally, we demonstrate that a Walrasian equilibrium may not exist, when the underlying value graph is more general. However, we also establish that in this case a more general pricing equilibrium always exists, and provide a stronger linear programming formulation that can be used to identify the efficient allocation for general graphical valuations.

We then consider solutions of these linear programming formulations using iterative algorithms. Complementing these iterative algorithms with appropriate payment rules, we obtain iterative auction formats that implement the efficient outcome at an (ex-post perfect) equilibrium. The auction formats we obtain rely on simple pricing rules that, in the most general case, require offering a bidder-specific price for each item, and bidder-specific discounts/markups for pairs of items. Our results in this part of the thesis suggest that when value functions of bidders exhibit some special structure, it is possible to systematically exploit this structure in order to develop simple efficient iterative auction formats.

Thesis Supervisor: Asuman Ozdaglar
Title: Professor

Thesis Supervisor: Pablo A. Parrilo
Title: Professor

To my families: past, present, and future.

Acknowledgments

Being a doctoral student at one of the most intellectually stimulating schools in the world has been a privilege. However, for me perhaps the most exciting and gratifying component of my MIT experience has been the opportunity to work with my advisors Asu Ozdaglar and Pablo Parrilo. Their guidance, combined with their endless support and enthusiasm made this thesis possible. Additionally, the nurturing and friendly research environment they created for their students at MIT, which is filled with innate curiosity and interesting ideas, has been essential for my development as a scholar and a researcher. I, with no doubt, have been very fortunate to complete my doctoral studies under their supervision. Both Asu and Pablo are outstanding examples for me, and I will try my best to carry their inspiration and guidance throughout my life.

I am also grateful to Prof. Georgia Perakis, and Prof. Daron Acemoglu for serving on my thesis committee. Their insights and perspective inspired me to refine and improve the ideas that are central to thesis, and discover new applications for my results. Consequently, this thesis has benefited enormously from their comments and feedback. I have also been very lucky to have their guidance and support in the job market. I would like to thank them for their dedication and help.

I am indebted to faculty members at LIDS, ORC, and CSAIL, especially to John Tsitsiklis, Devavrat Shah, Munther Dahleh, Itai Ashlagi, Vivek Farias, and Constantinos Daskalakis for creating a diverse and rich learning environment at MIT. Their lectures, ideas, and talks have been both inspiring and stimulating.

I have learned vastly from my other coauthors as well: Kostas Bimpikis, Ilan Lobel, Hamid Nazerzadeh, Jennifer Chayes, Christian Borgs, Ishai Menache, Constantinos Daskalakis, and Christos Papadimitriou. I have been extremely fortunate to work with some of these brightest minds from different communities. I look forward to interacting more with them, and having many other fruitful collaborations. I also want to thank Kostas, Ilan, and Hamid, for being great friends who are always willing to listen.

I am deeply grateful to (past and present) members of the Microsoft Research lab in New England, in particular, Jennifer Chayes, Christian Borgs, Ishai Menache, Adam Kalai, Sham Kakade, Brendan Lucier, and Yash Kanoria. In addition to the excellent talks and workshops I attended at Microsoft Research, my interaction with the members of the lab during my internships enriched my experience as a scholar, and exposed me to different ideas and perspectives. Also a special thanks goes to Jennifer for her support and help during my job search.

Various friends from ORC and LIDS made MIT a second home for me. I especially would like to thank Alireza Tahbaz-Salehi, Ercan Yildiz, Mohamed Mostagir, Azarakhsh Malekian, Christina Lee, Diego Feijer, Ali ParandehGheibi, Kuang Xu, Ammar Ammar, Spyros Zoumpoulis, Mihalis Markakis, Noah Stein, Amir Ali Ahmadi, Parikshit Shah, Paul Njoroge, James Saunderson, Takashi Tanaka, Dan Iancu, Mitra Osqui, Jagdish Ramakrish-

nan, Hamza Fawzi, Venkat Chandrasekaran, Srikanth Jagabathula, Hoda Eydgahi, Joline Uichanco, Ruben Lobel, Yehua Wei, Fernanda Bravo, Chaitanya Bandi, Gerry Tsoukalas, for making my years at MIT memorable.

Of course, my officemates Ermin Wei, Kimon Drakopoulos, Elie Adam, Annie Chen, and Jenny Lee have a special place in my LIDS family. I thank them for many fruitful discussions, their guidance, and support. But more importantly I am indebted to them for their friendship. I never imagined that a workplace can be this much fun!

I am also lucky to have many other great friends in Boston. I especially would like to thank Sefa Demirtas, Selda Celen, Ilke Kalcioglu, Alp Simsek, Sertac Karaman, Aylin Kentkur, Ozgur Amac, Guner Celik, Hakan Sonmez, Yalcin Cayir, Eray Sabancilar, Halil Tekin, Orcun Kurugol, for many fun weekends, trips to beaches and ski resorts, cookouts, and potlucks.

I am deeply indebted to my dear girlfriend Alice Fan for her endless kindness and love. Her support and companionship, not only helped me tackle the obstacles during the difficult times of my doctoral studies, but also brought me joy and changed my life in many different ways. I have been very fortunate to have her in my life. I owe her much more than I would ever be able to express.

Finally, I would like to extend my deepest gratitude to my parents Siddika and Zafer, and brother Utkan. I owe the opportunities and success I have had thus far to their sacrifices and guidance. Their unconditional love, support, and patience enabled me to successfully complete my doctoral studies.

Contents

1	Introduction	13
1.1	Main Contributions and Outline	17
1.1.1	Part I: Structural Game Decompositions and Dynamics	17
1.1.2	Part II: Iterative Auction Design for Graphical Valuations	18
I	Structural Game Decompositions and Dynamics	23
2	Decomposition of Games	25
2.1	Introduction	25
2.1.1	Related Literature	29
2.1.2	Outline	30
2.2	Game-Theoretic Background	31
2.2.1	Preliminaries	31
2.2.2	Games and Flows on Graphs	32
2.3	Flows and Helmholtz Decomposition	34
2.4	Canonical Decomposition of Games	38
2.4.1	Preliminaries	39
2.4.2	Decomposition of Games	43
2.4.3	An Example: Decomposition of Bimatrix Games	48
2.5	Potential and Harmonic Games	49
2.5.1	Potential Games	50
2.5.2	Harmonic Games	51
2.5.3	Nonstrategic Component and Efficiency in Games	57
2.5.4	Zero-Sum Games and Identical Interest Games	60
2.6	Projections onto Potential and Harmonic Games	62
2.7	Summary	64
2.8	Appendix: Additional Proofs	64
2.9	Appendix: Equilibria in Two Player Harmonic Games	70

3	Dynamics in Near-Potential Games	75
3.1	Introduction	75
3.1.1	Related Literature	80
3.1.2	Outline	82
3.2	Preliminaries	82
3.3	Better Response and Best Response Dynamics	84
3.4	Logit Response Dynamics	87
3.4.1	Properties of Logit Response	87
3.4.2	Stationary Distribution of Logit Response Dynamics	89
3.5	Fictitious Play	95
3.5.1	Mixed Strategies and Equilibria	95
3.5.2	Discrete-Time Fictitious Play	97
3.6	Summary	106
3.7	Appendix: Proofs of Section 3.5	106
II	Iterative Auction Design for Graphical Valuations	117
4	Tree Valuations	119
4.1	Introduction and Background	119
4.1.1	Related literature	122
4.1.2	Outline	124
4.2	Model and Preliminaries	125
4.2.1	Graphical Valuations	125
4.2.2	Anonymous Item Pricing and Walrasian Equilibrium	129
4.3	Walrasian Equilibrium and Tree Valuations	131
4.3.1	An Alternative LP Formulation	132
4.3.2	Existence of a Walrasian Equilibrium	137
4.4	Relaxing the Assumptions	138
4.4.1	Sign Consistency	138
4.4.2	Value Graphs with Cycles	139
4.5	Summary	141
4.6	Appendix: Gross Substitutes and Complements Condition and Graphical Valuations	143
4.7	Appendix: Proofs and Additional Results	146
5	General Graphical Valuations	157
5.1	Introduction and Organization	157
5.2	Anonymous Graphical Pricing	160
5.3	Bidder-Specific Item Pricing	164

5.4	Bidder-Specific Graphical Pricing	167
5.5	Additively Decomposable Valuations	169
5.6	Summary	171
5.7	Appendix: Proof of Theorem 5.2.1	172
6	Iterative Algorithms and Auction Design	179
6.1	Introduction and Organization	179
6.2	Iterative Algorithms for Linear Programs	182
6.3	Ex-post Perfect Equilibrium	184
6.4	Iterative Auctions for Tree Valuations	188
6.4.1	An Iterative Algorithm for LP1	188
6.4.2	An Efficient Iterative Auction	193
6.5	Iterative Auctions for General Graphical Valuations	197
6.6	Generalization: An Alternative LP and Additively Decomposable Valuations	203
6.7	Summary	208
6.8	Appendix: Additional Proofs	209
7	Conclusions	221
7.1	Summary	221
7.2	Future Research: Decomposition & Dynamics	222
7.3	Future Research: Graphical Valuations and Mechanism Design	223

Chapter 1

Introduction

Traditional systems engineering assumes the presence of a central entity, who has full control over the system, and designs algorithms that optimize a single, system-wide objective. This central entity usually has a well-defined control objective, and access to sufficient information to optimize for that objective (Khalil, 2002; Bertsekas and Tsitsiklis, 1989).

Modern systems, in contrast, involve self-interested agents with a diverse set of service requirements and potentially conflicting objectives. These agents take actions to maximize their objectives, without necessarily taking into account the effect of their actions on the remaining agents in the system. Examples include institutions in financial networks, firms in markets, and engineering systems with autonomous agents (such as communication networks, online computing infrastructures, electric power systems, robotic systems, and traffic networks). The presence of self-interested agents in these systems necessitate *the analysis of* the impact of various strategic considerations on the overall system, and *the design of* new systems with improved performance guarantees.

Analysis of strategic interactions, however, can be an intractable task in general game-theoretic settings. For instance, conceptually basic problems such as characterizing the Nash equilibria of finite games in strategic form, turn out to be intractable from a computational point of view unless the underlying strategic interactions exhibit some special properties (Daskalakis et al., 2006; Daskalakis and Papadimitriou, 2005; Nisan et al., 2007). Analogously, characterizing the outcome of dynamic strategic interactions is a difficult problem even for two player games, and there is no systematic framework for analyzing the limiting behavior of many of the adaptive update rules in general game-theoretic settings (Jordan, 1993; Fudenberg and Levine, 1998; Shapley, 1964).

Similar difficulties also arise in the design of mechanisms that maximize a global objective in the presence of strategic agents. Consequently, potentially inefficient/suboptimal mechanisms are employed in large scale social and economic systems. For instance, such multi-item auction mechanisms are used by governments for selling spectrum bands, by regional transmission organizations (such as New England ISO) to purchase sufficient electricity capacity in the forward capacity market, and by various businesses and govern-

ments to procure goods and services. One explanation for the use of such potentially inefficient/suboptimal mechanisms in practice is the difficulty of implementing the efficient/optimal outcome in social and economic systems involving a large number of agents, who have complex and potentially conflicting preferences, and interact repeatedly in a networked dynamic environment. Such features of social and economic systems may, at full generality, render the problem of implementing a desirable outcome in these systems intractable (Nisan and Segal, 2006; Nisan et al., 2007).

These observations motivate two exciting and complementary research directions that are the key components of this thesis. First, it is imperative to conduct fundamental research to *develop improved models and theoretical tools for analyzing multi-agent interactions* in various social and economic systems, and understanding their main features. Second, it is necessary to *develop simpler and better mechanisms* with improved efficiency and optimality guarantees, by exploiting these features.

Even though analyzing the outcome of strategic interactions of a large number of agents who have complex preferences, is a challenging task in general, it is known that special classes of games exhibit desirable equilibrium and dynamic properties, and allow for tractable analysis. An important class of games with well-understood equilibrium and dynamic properties is potential games (see e.g. Young (2004); Monderer and Shapley (1996b); Marden et al. (2009b)). These games have potential functions, which summarize incentives of different players jointly. The limiting behavior of various update rules including better/best response dynamics (Monderer and Shapley, 1996b; Young, 2004), fictitious play (Monderer and Shapley, 1996a; Shamma and Arslan, 2004; Marden et al., 2009b; Hofbauer and Sandholm, 2002) and logit response dynamics (Blume, 1993, 1997; Alós-Ferrer and Netzer, 2010; Marden and Shamma, 2012) can be established for potential games using the corresponding potential function.

Motivated by this observation, in the first part of this thesis, we investigate why classes of games such as potential games have desirable static and dynamic properties. In particular, we focus on structural properties of potential games, and delineate the fundamental characteristics of preferences of players, which lead to these properties. Additionally, by introducing an alternative flow representation for finite games in strategic form, and employing tools from algebraic topology, we develop a canonical direct sum decomposition of an arbitrary game into three components, which we refer to as the *potential*, *harmonic* and *nonstrategic* components. The first component of this decomposition captures the desirable strategic properties of potential games, whereas the second component leads to qualitatively different equilibrium and dynamic properties. By exploiting this decomposition, we develop a novel approach for approximating a given game using a related potential game.

Additionally, we show that this decomposition provides a valuable tool for the analysis of equilibria and dynamics in games. In particular, we first establish that the set of approximate equilibria of an arbitrary game can be characterized through the equilibria of

its potential game approximation. We then extend our analysis to the characterization of dynamics in games, and provide a framework that can be used to characterize the limiting behavior of dynamic strategic interactions in a given game, in terms of the outcome of dynamics in its potential game approximation. The results of this part of the thesis provide new theoretical tools for the (approximate) characterization of equilibria and dynamics in various game-theoretic settings.

In the second part of the thesis, we change our focus to implementing the efficient outcomes in multi-agent settings through simple mechanisms. In particular, we develop novel iterative auction formats that guarantee efficiency in environments where the auctioneer sells multiple items that can exhibit value complementarities and substitutabilities.

Iterative auctions are a class of mechanisms that are commonly employed in practice. In these auctions, the auctioneer sets prices for the items she is selling, bidders report which items they are interested in at the given prices, and in response to these reports, the auctioneer updates the prices. The well-known English and Dutch auctions are examples of single-item iterative auctions. When bidders have independent private values, these auctions allocate the item efficiently, i.e., the bidder with the highest value receives the item (Krishna, 2009).

Arguably, iterative auction formats are more common in practice than their static counterparts (such as sealed bid auctions), due to their desirable properties such as privacy preservation, price discovery, and reduced communication requirements (Ausubel and Milgrom, 2006; Rothkopf et al., 1990; Engelbrecht-Wiggans and Kahn, 1991). For this reason, a number of papers in the recent literature focused on the question of designing efficient iterative multi-item auctions. Examples include, package bidding auction (Ausubel and Milgrom, 2002), clinching auction and its variants (Ausubel, 2004, 2006), and auctions that rely on universally competitive equilibria (UCE) (Mishra and Parkes, 2007). Other examples, which focus explicitly on myopic strategy updates, include best response mechanisms of (Nisan et al., 2011b) and (Nisan et al., 2011a).

In general multi-item settings (such as spectrum or procurement auctions) the iterative auction formats that are present in literature do not always guarantee efficiency. More precisely, either they guarantee efficiency under some restrictive assumptions (such as the gross substitutes assumption, Gul and Stacchetti (2000); Ausubel (2006)), or they rely on complex pricing rules that require offering a different price for each bundle of items the auctioneer sells (Bikhchandani et al., 2002; De Vries et al., 2007; Ausubel and Milgrom, 2002; Ausubel, 2006; Mishra and Parkes, 2007; Vohra, 2011). The auction formats in the first category do not allow for value complementarity between different items, which is commonly observed in practical auction environments. Those in the second category, on the other hand, may not be practical. This is because these auctions require reporting exponentially many prices to the bidders at each stage of the auction.

These observations motivate us to design novel iterative auctions for multi-item environ-

ments. Our main contribution in the second part of the thesis is to develop simple efficient iterative auction formats for settings that involve both complementarity and substitutability in valuations. We obtain our results by focusing on a special class of valuation functions, which we refer to as graphical valuations. These valuations are not fully general, but importantly they allow for a compact representation of the value functions of the bidders, and capture the structural properties of valuations in important combinatorial auctions. Due to this compact structure, unlike the auctions in the existing literature, the iterative auction formats we provide in this part of the thesis rely on simple pricing rules and guarantee efficiency even in settings where valuation functions exhibit both value complementarity and substitutability.

Our approach for developing iterative auctions involves three steps. First, we focus on providing linear programming formulations of the efficient allocation problem. Then, we consider iterative algorithms for solutions of these formulations, and show that these suggest a natural price/demand update process that converges to the efficient outcome. Finally, we obtain iterative auctions by complementing these algorithms with appropriate final payment schemes. The final step also guarantees that in the iterative auctions we develop, it is an equilibrium for bidders to reveal their demand truthfully, and this equilibrium leads to an efficient allocation.

This approach for iterative auction design is also employed in the existing literature (see Vohra (2011)). However, the existing iterative auction formats that follow this approach and allow for complementarity in valuations, rely on exponentially many prices for implementing the efficient outcome (Bikhchandani et al., 2002; De Vries et al., 2007; Mishra and Parkes, 2007; Vohra, 2011). In contrast, our main contribution in this part of the thesis is to develop efficient iterative auction formats that rely on simple pricing rules. We accomplish this by following the approach outlined above, and carefully exploiting the structural properties of graphical valuations, in order to obtain simple optimization formulations of the efficient allocation problem, and ultimately iterative auction formats that rely on simple pricing rules.

Our results in this part of the thesis suggest that when valuation functions of bidders exhibit some special structure, it is possible to systematically exploit this structure, in order to develop simple efficient iterative auction formats. These iterative auctions rely on pricing rules that have a similar structure to that of the underlying valuations. Hence, it is possible to implement the efficient outcome using a pricing rule that is no more complex than the valuation functions. Therefore, by first identifying the structure in valuations of bidders, and then following the framework we propose in this thesis for iterative auction design, it may be possible to obtain simple iterative auction formats that are applicable in practice.

A more detailed summary of the main contributions of each of the chapters of this thesis, and an outline are provided in the next section.

1.1 Main Contributions and Outline

The remainder of this thesis is divided into two parts. The first part contains our contributions on structural game decompositions (Chapter 2) and analysis of dynamic strategic interactions (Chapter 3). The second part focuses on the question of iterative auction design for graphical valuations. In this part, we first obtain a linear programming formulation of the efficient allocation problem for a subclass of graphical valuations, where the underlying graph has a tree structure, i.e., it does not have any cycles (Chapter 4). However, we also show that this formulation may not be used to find the efficient outcome, if the underlying graph is not a tree. In Chapter 5 we provide alternative linear programming formulations that can be solved to identify the efficient outcome for general graphical valuations. In Chapter 6, we use the solutions of these linear programs with iterative algorithms to develop iterative auction formats. Additionally, we characterize the equilibria of our auctions in Chapter 6, and establish that the efficient outcome can be implemented at an ex-post perfect equilibrium by complementing the auctions with appropriate payment schemes. Future directions related to the contributions of this thesis are outlined in Chapter 7.

We conclude this chapter with a detailed summary of the main contributions of the remaining chapters.

1.1.1 Part I: Structural Game Decompositions and Dynamics

Chapter 2. In this chapter we introduce a novel flow representation for finite games in strategic form. This representation allows us to develop a canonical direct sum decomposition of an arbitrary game into three components, which we refer to as the *potential*, *harmonic*, and *nonstrategic* components. We analyze natural classes of games that are induced by this decomposition, and in particular, focus on games with no harmonic component and games with no potential component. We show that the first class corresponds to the well-known *potential games*. We refer to the second class of games as *harmonic games*, and study the structural and equilibrium properties of this new class of games.

Intuitively, the potential component of a game captures interactions that can equivalently be represented as a common interest game, while the harmonic part represents the conflicts between the interests of the players. We make this intuition precise, by studying the properties of these two classes, and show that indeed they have quite distinct and remarkable characteristics. For instance, while finite potential games always have pure Nash equilibria, harmonic games generically never do. Moreover, we show that the nonstrategic component does not affect the equilibria of a game, but plays a fundamental role in their efficiency properties, thus decoupling the location of equilibria and their payoff-related properties. Exploiting the properties of the decomposition framework, we obtain explicit expressions for the projections of games onto the subspaces of potential and harmonic games. This enables an extension of the properties of potential and harmonic games to “nearby” games.

We exemplify this point by showing that the set of approximate equilibria of an arbitrary game can be characterized through the equilibria of its projection onto the set of potential games.

Chapter 3. Except for special classes of games, there is no systematic framework for analyzing the dynamical properties of multi-agent strategic interactions. Potential games are one such special but restrictive class of games that allow for tractable analysis of dynamics. Intuitively, games that are “close” to a potential game should share similar properties. In this chapter, we formalize and develop this idea by quantifying to what extent the dynamic features of potential games extend to “near-potential” games.

We study convergence of three commonly studied classes of adaptive dynamics: discrete-time better/best response, logit response, and discrete-time fictitious play dynamics. For better/best response dynamics, we focus on the evolution of the sequence of pure strategy profiles and show that this sequence converges to a (pure) approximate equilibrium set, whose size is a function of the “distance” from a close potential game. We then study logit response dynamics parametrized by a smoothing parameter that determines the frequency with which the best response strategy is played. Our analysis uses a Markov chain representation for the evolution of pure strategy profiles. We provide a characterization of the stationary distribution of this Markov chain in terms of the distance of the game from a close potential game and the corresponding potential function. We further show that the stochastically stable strategy profiles (defined as those that have positive probability under the stationary distribution in the limit as the smoothing parameter goes to 0) are pure approximate equilibria. Finally, we turn our attention to fictitious play, and establish that in near-potential games, the sequence of empirical frequencies of player actions converges to a neighborhood of (mixed) equilibria of the game, where the size of the neighborhood increases with the distance of the game to a potential game. Thus, our results suggest that games that are close to a potential game inherit the dynamical properties of potential games. Since a close potential game to a given game can be found using the game decomposition results, our approach also provides a systematic framework for studying convergence behavior of adaptive learning dynamics in arbitrary finite strategic form games.

1.1.2 Part II: Iterative Auction Design for Graphical Valuations

Chapter 4. We start this chapter by introducing graphical valuations. Value functions that belong to this class are associated with a value graph, nodes of which correspond to the items that are sold by the auctioneer. The edges of this graph capture the value complementarity and substitutability exhibited by the items. The value a bidder has for a given set of items can be expressed as the sum of the weights of nodes and edges that are contained in this set. This valuation model captures the value complementarity/substitutability structure in important practical settings, while allowing for a compact representation of the

value functions. A closely related valuation model appeared in a recent work by (Abraham et al., 2012), where the focus was on the complexity of auction design for (hyper) graphical valuations that do not exhibit substitutabilities. In contrast, in this work we develop simple efficient iterative auctions for graphical valuations that exhibit both value complementarities and substitutabilities.

An important component of iterative auction design is the choice of the pricing rule used for running auctions. In this chapter, we present an important pricing rule, anonymous item pricing, that is commonly used in the literature for the design of iterative auctions (Ausubel, 2006). We also discuss a natural termination condition for iterative auctions that rely on this pricing rule: the auctioneer terminates the auction when a “market clearance” condition holds, i.e., when all bidders demand disjoint sets of items, and all items are demanded by some bidder. It is clear that at such an outcome no bidder needs compete with the remaining bidders to acquire the set of items that she demands (since the demand sets are disjoint), thereby making this outcome a natural termination point for the auction. Moreover, this termination condition is equivalent to convergence of the iterative auctions to a Walrasian equilibrium. Hence, it is possible to design iterative auction formats that rely on anonymous item pricing and the aforementioned termination condition if and only if a Walrasian equilibrium exists.

For the remainder of this chapter, we restrict attention to a special subclass of graphical valuations, where the underlying value graph is a tree, and valuations satisfy an additional technical (sign consistency) condition. For this class of valuation functions, we show that a Walrasian equilibrium always exists. It is known that the existence of a Walrasian equilibrium is equivalent to existence of integral optimal solutions to a linear programming formulation of the efficient allocation problem. Thus, our result immediately leads to a linear program that can be solved to identify the efficient allocation for sign-consistent tree valuations. We also demonstrate that if we relax the sign consistency assumption, or the tree assumption, solving this linear programming formulation no longer gives the efficient outcome and a Walrasian equilibrium does not exist.

Chapter 5. In this chapter, we study the efficient allocation problem for more general value graphs. As established in the previous chapter, for general graphical valuations, a Walrasian equilibrium need not exist, hence it is not possible to develop iterative auctions that terminate at a market clearing outcome using anonymous item pricing. This motivates us to consider more general pricing rules, and iterative auction formats that terminate when a generalized market clearance condition with such pricing rules holds. To this end, we first introduce the concept of a pricing equilibrium, which is a generalization of the Walrasian equilibrium concept to pricing rules that are more general than anonymous item pricing. Then, we provide linear programming formulations of the efficient allocation problem, which have optimal solutions that are integral (hence identify the efficient outcome) if and only if

pricing equilibria exist. Iterative solutions of these LP formulations can be used to develop efficient iterative auction formats that terminate when a pricing equilibrium is found, as we discuss in detail in Chapter 6.

More precisely, in this chapter, we focus on three pricing rules that generalize anonymous item pricing, and the associated linear programming formulations. The first linear programming formulation we focus on strengthens the formulation of Chapter 4 by imposing a constraint for each edge of the underlying value graph. Hence, its dual associates an anonymous price variable with each node, and edge of the underlying value graph. We refer to this pricing rule as anonymous graphical pricing. Using this linear programming formulation and its dual, we establish that a pricing equilibrium with anonymous graphical pricing exists if and only if this formulation has optimal solutions that are integral. We also establish that when the underlying value graph involves a 5-clique (as a minor), then it may be possible to find the efficient allocation using this linear programming formulation even in cases where the formulation of Chapter 4 does not obtain the efficient allocation. Conversely, if the value graph does not involve a 4-clique (as a minor), then the two formulations are equivalent, in the sense that if one formulation gives the efficient outcome so does the other one, and vice versa. These results suggest that even stronger optimization formulations (or more general pricing rules) may be necessary in order to find the efficient outcome for more general value graphs.

We then study linear programming formulations, whose duals suggest bidder-specific item pricing, and bidder-specific graphical pricing rules. These pricing rules are analogous to the anonymous pricing rules mentioned before, but they allow for offering different prices to different bidders. We show that the first linear programming formulation is equivalent to that of Chapter 4, and hence cannot find the efficient allocation for general value graphs. On the other hand, the formulation associated with bidder-specific graphical pricing is stronger than all of the aforementioned formulations, and can be used to identify the efficient outcome for all graphical valuations. This result also implies that a pricing equilibrium with this pricing rule always exists. Moreover, this formulation can be generalized to obtain linear programming formulations (and pricing equilibria) that identify the efficient outcome, even for valuation functions that exhibit a more general additively decomposable structure than graphical valuations.

Chapter 6. In this chapter, we first introduce the solution concept, ex-post perfect equilibrium, which we use for the analysis of iterative auctions, and provide conditions for characterization of such equilibria. Then we obtain iterative algorithms for the solution of the linear programming formulations of Chapters 4 and 5, and employ these algorithms to develop iterative auction formats for graphical valuations. These auctions imitate the iterative algorithms if bidders reveal their demand truthfully. Moreover, by charging bidders appropriate final payments, we show that it is an ex-post perfect equilibrium for bidders

to truthfully reveal their demand in these auctions. These results imply that our iterative auctions guarantee efficiency at an ex-post perfect equilibrium.

More precisely, we first focus on the linear programming formulation of Chapter 4 that allows for finding the efficient outcome for sign-consistent tree valuations. We show that an iterative solution of this formulation using primal-dual algorithms suggests a natural iterative auction format, where the auctioneer offers a single anonymous price for each item. At each stage of this auction, the auctioneer increases the prices of the overdemanded items, and decreases those of the underdemanded ones. This auction terminates when a market clearance condition holds and allocates items efficiently, if bidders reveal their demand truthfully. However, bidders may have incentive to misreport their demand, if their final payments are equal to the prices that emerge at the end of the auction. We show that by appropriately modifying these payments, the auctioneer can guarantee truthful bidding at each stage of the auction. The corresponding iterative auction format guarantees efficiency at an (ex-post perfect) equilibrium for sign-consistent tree valuations.

We then follow a similar approach for developing iterative auctions that guarantee efficiency for general graphical valuations. In particular, we focus on iterative algorithms that can be used for solving the linear programming formulation of Chapter 5 associated with bidder-specific graphical pricing. These iterative algorithms can be used to find the efficient allocation for all graphical valuations, since the corresponding linear programming formulation has an optimal solution associated with such allocations. Additionally, we establish that the aforementioned linear program has an optimal solution that allows for computing final payments that guarantee truthful bidding by bidders. Employing these payments together with our iterative algorithm, we provide an iterative auction format that (terminates at a pricing equilibrium, and) implements the efficient outcome at an ex-post perfect equilibrium.

In this chapter, we also discuss how our results and auction formats can be generalized to environments where valuations of bidders are not necessarily graphical, but admit a more general additively decomposable structure. The results of this part of the thesis imply that in general it is possible to develop efficient iterative auction formats that rely on pricing rules that have a similar structure to the underlying valuation functions. Hence, in practice it may be possible to develop simple iterative auction formats, by first identifying the structure of the valuations of bidders, and then following the framework provided in this part of the thesis to exploit this special structure.

Part I

Structural Game Decompositions and Dynamics

Chapter 2

Decomposition of Games

2.1 Introduction

Potential games play an important role in game-theoretic analysis due to their desirable static properties (e.g., existence of a pure strategy Nash equilibrium) and tractable dynamics (e.g., convergence of simple player dynamics to a Nash equilibrium); see Monderer and Shapley (1996b,a), and Neyman (1997). However, many multi-agent strategic interactions in economics and engineering cannot be modeled as a potential game.

This chapter provides a novel flow representation of the preference structure in strategic-form finite games, which allows for delineating the fundamental characteristics in preferences that lead to potential games. This representation enables us to develop a canonical orthogonal decomposition of an arbitrary game into a potential component, a harmonic component, and a nonstrategic component, each with its distinct properties. The decomposition can be used to define the “distance” of an arbitrary game to the set of potential games. We use this fact to describe the approximate equilibria of a given game in terms of the equilibria of its potential component. Moreover, we illustrate in the next chapter that a similar approach allows for characterizing the outcome of dynamic strategic interactions in games.

Our starting point is to associate to a given finite game a *game graph*, where the set of nodes corresponds to the strategy profiles and the edges represent the “comparable strategy profiles” i.e., strategy profiles that differ in the strategy of a single player. The utility differences for the deviating players along the edges define a flow on the game graph. Although this graph contains strictly less information than the original description of the game in terms of utility functions, all relevant strategic aspects (e.g., equilibria) are captured.

Our first result provides a *canonical decomposition* of an arbitrary game using tools from the study of flows on graphs (which can be viewed as combinatorial analogues of vector fields). In particular, we use the *Helmholtz decomposition theorem* (e.g., Jiang et al. (2011)), which enables the decomposition of a flow on a graph into three components: globally consistent, locally consistent (but globally inconsistent), and locally inconsistent component (see Theorem 2.3.1). The globally consistent component represents a gradient flow while the

locally consistent flow corresponds to flows around global cycles. The locally inconsistent component represents local cycles (or circulations) around 3-cliques of the graph.

Our game decomposition has three components: *nonstrategic*, *potential* and *harmonic*. The first component represents the “nonstrategic interactions” in a game. Consider two games in which, given the strategies of the other players, each player’s utility function differs by an additive constant. These two games have the same utility differences, and therefore they have the same flow representation. Moreover, since equilibria are defined in terms of utility differences, the two games have the same equilibrium set. We refer to such games as *strategically equivalent*. We normalize the utilities (by adding constants to the utilities of a player given the other players’ strategies such that their sum is equal to zero), and refer to the utility differences between a game and its normalization as the nonstrategic component of the game. Our next step is to remove the nonstrategic component and apply the Helmholtz decomposition to the remainder. The flow representation of a game defined in terms of utility functions (as opposed to preferences) does not exhibit local cycles, therefore the Helmholtz decomposition yields the two remaining components of a game: the *potential component* (gradient flow) and the *harmonic component* (global cycles). The decomposition result is particularly insightful for bimatrix games (i.e., finite games with two players, see Section 2.4.3), where the potential component represents the “team part” of the utilities (suitably perturbed to capture the utility matrix differences), and the harmonic component corresponds to a zero-sum game.

The canonical decomposition we introduce is illustrated in the following example.

Example 2.1.1 (Road-sharing game). *Consider a three-player game, where each player has to choose one of the two roads $\{0, 1\}$. We denote the players by d_1 , d_2 and s . The player s tries to avoid sharing the road with other players: its payoff decreases by 2 with each player d_1 and d_2 who shares the same road with it. The player d_1 receives a payoff -1 , if d_2 shares the road with it and 0 otherwise. The payoff of d_2 is equal to negative of the payoff of d_1 , i.e., $u^{d_1} + u^{d_2} = 0$. Intuitively, player d_1 tries to avoid player d_2 , whereas player d_2 wants to use the same road with d_1 .*

In Figure 2-1a we present the flow representation for this game (described in detail in Section 2.2.2), where the nonstrategic component has been removed. Figures 2-1b and 2-1c show the decomposition of this flow into its potential and harmonic components. In the figure, each tuple (a, b, c) denotes a strategy profile, where player s uses strategy a and players d_1 and d_2 use strategies b and c respectively.

This example shows that the harmonic component of a game satisfies the flow conservation condition, i.e., the total flow at each node is equal to zero. On the other hand, for the potential component, the total flow around every cycle is equal to zero. These observations highlight a key distinction in the flow representations of harmonic and potential components of a game: the harmonic component can be characterized by the presence of “preference cycles”, while the potential component does not have such cycles.

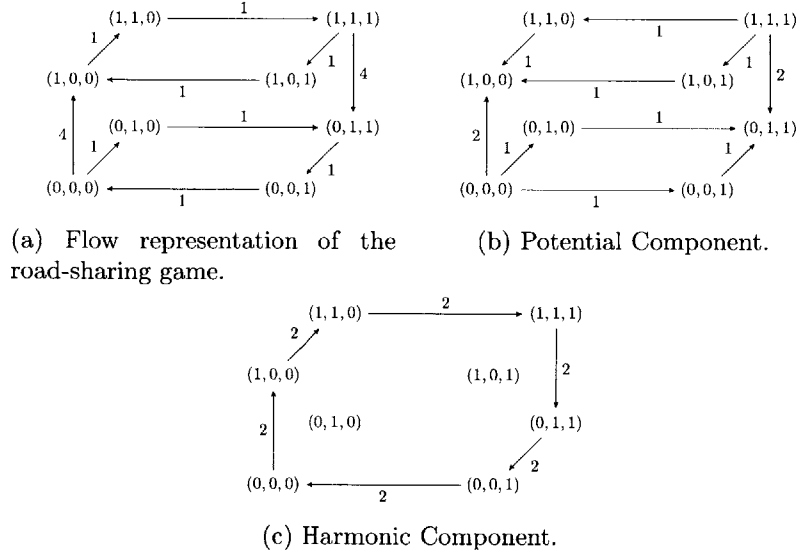


Figure 2-1: Potential-harmonic decomposition of the road-sharing game. An arrow between two strategy profiles, indicates the improvement direction in the payoff of the player who changes its strategy, and the associated number quantifies the improvement in its payoff.

These components induce a direct sum decomposition of the space of games into three respective subspaces, which we refer to as the *nonstrategic*, *potential* and *harmonic* subspaces, denoted by \mathcal{N} , \mathcal{P} , and \mathcal{H} , respectively. We use these subspaces to define classes of games with distinct equilibrium properties. We establish that the set of potential games coincides with the direct sum of the subspaces \mathcal{P} and \mathcal{N} , i.e., potential games are those with no harmonic component. Similarly, we define a new class of games in which the potential component vanishes as *harmonic games*. The classical rock-paper-scissors and matching pennies games are examples of harmonic games. The decomposition then has the following structure:

$$\underbrace{\mathcal{P}}_{\text{Potential games}} \oplus \underbrace{\mathcal{N} \oplus \mathcal{H}}_{\text{Harmonic games}}.$$

It is insightful to provide alternative definitions of potential and harmonic games in terms of the payoff functions of players. A game with a set of players \mathcal{M} , set of strategies E^m for all $m \in \mathcal{M}$, and a collection of utility functions $\{u^m\}_{m \in \mathcal{M}}$ is a *potential game* if there exists a potential function ϕ satisfying

$$\phi(\mathbf{p}^m, \mathbf{p}^{-m}) - \phi(\mathbf{q}^m, \mathbf{p}^{-m}) = u^m(\mathbf{p}^m, \mathbf{p}^{-m}) - u^m(\mathbf{q}^m, \mathbf{p}^{-m}), \quad (2.1)$$

for every $m \in \mathcal{M}$, $\mathbf{p}^m, \mathbf{q}^m \in E^m$, $\mathbf{p}^{-m} \in E^{-m}$. The game is a *harmonic game* if for all

strategy profiles \mathbf{p} and strategies $\mathbf{q}^m \in E^m$, the utility functions satisfy

$$\sum_{m \in \mathcal{M}} \sum_{\mathbf{q}^m \in E^m} (u^m(\mathbf{p}^m, \mathbf{p}^{-m}) - u^m(\mathbf{q}^m, \mathbf{p}^{-m})) = 0. \quad (2.2)$$

It can be seen from these definitions that harmonic games are games which satisfy the flow conservation condition illustrated in Example 2.1.1, and potential games are those for which the total flow around each cycle is equal to zero. The equivalences of subspace definitions of potential and harmonic games, and the definitions in (2.1) and (2.2) are established in Sections 2.5.1 and 2.5.2.

Our second set of results establishes properties of potential and harmonic games and examines how the nonstrategic component of a game affects the efficiency of equilibria. Harmonic games can be characterized by the existence of improvement cycles, i.e., cycles in the game graph, where at each step the player that changes its action improves its payoffs. We show that harmonic games generically do not have pure Nash equilibria. Interestingly, for the special case when the number of strategies of each player is the same, a harmonic game satisfies a “multi-player zero-sum property” (i.e., the sum of utilities of all players is equal to zero at all strategy profiles). We also study the mixed Nash and correlated equilibria of harmonic games. We show that the uniformly mixed strategy profile (see Definition 2.5.2) is always a mixed Nash equilibrium and if there are two players in the game, the set of mixed Nash equilibria generically coincides with the set of correlated equilibria. Finally, we focus on the nonstrategic component of a game. As discussed above, the nonstrategic component does not affect the equilibrium set. Using this property, we show that by changing the nonstrategic component of a game, it is possible to make the set of Nash equilibria coincide with the set of Pareto optimal strategy profiles. Thus, while this component does not change the equilibrium set, it determines the efficiency properties of the equilibria.

Our third set of results focuses on the *projection of a game onto its respective components*. We first define a natural inner product and show that under this inner product the components in our decomposition are orthogonal. We further provide explicit expressions for the closest potential and harmonic games to a game with respect to the norm induced by the inner product. We use the distance of a game to its closest potential game to characterize the approximate equilibrium set in terms of the equilibria of the potential game.

The decomposition framework in this chapter leads to the identification of subspaces of games with distinct and tractable equilibrium properties. Understanding the structural properties of these subspaces and the classes of games they induce, provides new insights and tools for analyzing the equilibrium properties of general noncooperative games. Additionally, as explained in the next chapter, the decomposition framework serves as a valuable tool for analysis of dynamics in games.

2.1.1 Related Literature

Besides the works already mentioned, this chapter is also related to several papers in the cooperative and noncooperative game theory literature:

The idea of decomposing a game (using different approaches) into simpler games which admit more tractable equilibrium analysis appeared even in the early works in the cooperative game theory literature. In Von Neumann and Morgenstern (Von Neumann and Morgenstern, 1947), the authors propose to decompose games with large number of players into games with fewer players. In (Marinacci, 1996; Gilboa and Schmeidler, 1995; Shapley, 1953), a different approach is followed: the authors identify cooperative games through the games' value functions (see (Von Neumann and Morgenstern, 1947)) and obtain decompositions of the value function into simpler functions. By defining the component games using the simpler value functions, they obtain decompositions of games. In this approach, the set of players is not made smaller or larger by the decomposition but the component games have simpler structure. Another method for decomposing the space of cooperative games appeared in Kleinberg and Weiss (Kleinberg and Weiss, 1986, 1985). In these papers, the algebraic properties of the space of games and the properties of the nullspace of the Shapley value operator (see Shapley (Shapley, 1953)) and its orthogonal complement are exploited to decompose games. This approach does not necessarily simplify the analysis of games but it leads to an alternative expression for the Shapley value (Kleinberg and Weiss, 1985). Our work is on decomposition of noncooperative games, and different from the above references since we explicitly exploit the properties of noncooperative games in our framework.

In the context of noncooperative game theory, a decomposition for games in normal form appeared in Sandholm (Sandholm, 2010b). In this paper, the author defines 2^M component games that are characterized by their sets of active and passive players, and provides a decomposition of normal form games to these components. This decomposition is then used to identify potential games: the original game is a potential game if and only if in each component game the active players have identical payoff functions. We note that our decomposition approach is different than this work in the properties of the component games. In particular, using the global preference structure in games, our approach yields decomposition of games to three components with distinct equilibrium properties, and these properties can be exploited to gain insights about the static and dynamic features of the original game.

A different decomposition of noncooperative games appeared in (Başar and Ho, 1974; Kalai and Kalai, 2010). This decomposition relies on identifying zero-sum and identical interest components for a given game, and is used for dealing with cooperation-related issues that emerge in games with strategic players. It is fundamentally different than the decomposition we propose, since one of its components is always zero-sum, whereas this is not the case in our decomposition.

Related ideas of representing finite strategic form games as graphs previously appeared

in the literature to study different solution concepts in normal form games (Goemans et al., 2005; Christodoulou et al., 2006). In these references, the authors focus on the restriction of the game graph to best-reply paths and analyze the outcomes of games using this subgraph.

In our work, the graph representation of games and the flows defined on this graph lead to a natural equivalence relation. Related notions of strategic equivalence are employed in the game theory literature to generalize the desirable static and dynamic properties of games to their equivalence classes (Moulin and Vial, 1978; Rosenthal, 1974; Morris and Ui, 2004; Voorneveld, 2000; Germano, 2006; Hammond, 2005; Hofbauer and Hopkins, 2005; Kannan and Theobald, 2010; Mertens, 2004; Hillas and Kohlberg, 2002). Moulin and Vial (Moulin and Vial, 1978) refer to games which have the same better-response correspondence as equivalent games and study the equilibrium properties of games which are equivalent to zero-sum games. In (Hammond, 2005; Hofbauer and Hopkins, 2005), the dynamic and static properties of certain classes of bimatrix games are generalized to their equivalence classes. Using the best-response correspondence instead of the better-response correspondence, the papers (Rosenthal, 1974; Morris and Ui, 2004; Voorneveld, 2000) define different equivalence classes of games. We note that the notion of strategic equivalence used in our work implies some of the equivalence notions mentioned above. However, unlike these papers, our notion of strategic equivalence leads to a canonical decomposition of the space of games, which is then used to extend the desirable properties of potential games to “close” games that are not strategically equivalent.

Despite the fact that harmonic games were not defined in the literature before (and thus, the term “harmonic” does not appear explicitly as such), specific instances of harmonic games were studied in different contexts. Hofbauer and Schlag (2000) study dynamics in “cyclic games” and obtain results about a class of harmonic games which generalize the matching pennies game. A parametrized version of Dawkins’ battle of the sexes game, which is a harmonic game under certain conditions, is studied in Smith and Hofbauer (1987). Other examples of harmonic games have also appeared in the buyer/seller game of Friedman (1991) and the crime deterrence game of Cressman and Morrison (1998).

2.1.2 Outline

The remainder of this chapter is organized as follows. In Section 2.2, we present the relevant game theoretic background and provide a representation of games in terms of graph flows. In Section 2.3, we state the Helmholtz decomposition theorem which provides the means of decomposing a flow into orthogonal components. In Section 2.4, we use this machinery to obtain a canonical decomposition of the space of games. We introduce in Section 2.5 natural classes of games, namely potential and harmonic games, which are induced by this decomposition and describe the equilibrium properties thereof. In Section 2.6, we define an inner product for the space of games, under which the components of games turn out to be orthogonal. Using this inner product and our decomposition framework we propose

a method for projecting a given game to the spaces of potential and harmonic games. We then apply the projection to study the equilibrium properties of “near-potential” games. We close in Section 2.7 with a summary of the main contributions of this chapter. Some of the proofs of this chapter are delegated to Section 2.8. Additional properties of two player harmonic games are discussed in Section 2.9.

2.2 Game-Theoretic Background

In this section, we describe the required game-theoretic background. Notation and basic definitions are given in Section 2.2.1. In Section 2.2.2, we provide an alternative representation of games in terms of flows on graphs. This representation is used in the rest of this chapter to analyze finite games.

2.2.1 Preliminaries

A (noncooperative) *strategic-form finite game* consists of:

- A finite set of players, denoted $\mathcal{M} = \{1, \dots, M\}$.
- Strategy spaces: A finite set of strategies (or actions) E^m , for every $m \in \mathcal{M}$. The joint strategy space is denoted by $E = \prod_{m \in \mathcal{M}} E^m$.
- Utility functions: $u^m : E \rightarrow \mathbb{R}$, $m \in \mathcal{M}$.

A (strategic-form) game instance is accordingly given by the tuple $\langle \mathcal{M}, \{E^m\}_{m \in \mathcal{M}}, \{u^m\}_{m \in \mathcal{M}} \rangle$, which for notational convenience will often be abbreviated to $\langle \mathcal{M}, \{E^m\}, \{u^m\} \rangle$.

We use the notation $\mathbf{p}^m \in E^m$ for a strategy of player m . A collection of players’ strategies is given by $\mathbf{p} = \{\mathbf{p}^m\}_{m \in \mathcal{M}}$ and is referred to as a strategy profile. A collection of strategies for all players but the m -th one is denoted by $\mathbf{p}^{-m} \in E^{-m}$. We use $h_m = |E^m|$ for the cardinality of the strategy space of player m , and $|E| = \prod_{m=1}^M h_m$ for the overall cardinality of the strategy space. As an alternative representation, we shall sometimes enumerate the actions of the players, so that $E^m = \{1, \dots, h_m\}$.

The basic solution concept in a noncooperative game is that of a *Nash Equilibrium* (NE). A strategy profile $\mathbf{p} \triangleq \{\mathbf{p}^1, \dots, \mathbf{p}^M\}$ is an ϵ -equilibrium if

$$u^m(\mathbf{p}^m, \mathbf{p}^{-m}) \geq u^m(\mathbf{q}^m, \mathbf{p}^{-m}) - \epsilon \quad \text{for every } \mathbf{q}^m \in E^m \text{ and } m \in \mathcal{M}. \quad (2.3)$$

A (pure) Nash equilibrium¹ is an ϵ -equilibrium with $\epsilon = 0$.

The next lemma shows that the ϵ -equilibria of two games can be related in terms of the differences in utilities.

¹In strategic-form finite games, Nash equilibrium always exists in mixed strategies. However, pure Nash equilibria, defined here need not always exist.

Lemma 2.2.1. *Consider two games \mathcal{G} and $\hat{\mathcal{G}}$, which differ only in their utility functions, i.e., $\mathcal{G} = \langle \mathcal{M}, \{E^m\}, \{u^m\} \rangle$ and $\hat{\mathcal{G}} = \langle \mathcal{M}, \{E^m\}, \{\hat{u}^m\} \rangle$. Assume that $|u^m(\mathbf{p}) - \hat{u}^m(\mathbf{p})| \leq \epsilon_0$ for every $m \in \mathcal{M}$ and $\mathbf{p} \in E$. Then, every ϵ_1 -equilibrium of $\hat{\mathcal{G}}$ is an ϵ -equilibrium of \mathcal{G} for some $\epsilon \leq 2\epsilon_0 + \epsilon_1$ (and viceversa).*

Proof. Let \mathbf{p} be an ϵ_1 -equilibrium of $\hat{\mathcal{G}}$ and let $\mathbf{q} \in E$ be a strategy profile with $\mathbf{q}^k \neq \mathbf{p}^k$ for some $k \in \mathcal{M}$, and $\mathbf{q}^m = \mathbf{p}^m$ for every $m \in \mathcal{M} \setminus \{k\}$. Then,

$$u^k(\mathbf{q}) - u^k(\mathbf{p}) \leq u^k(\mathbf{q}) - u^k(\mathbf{p}) - (\hat{u}^k(\mathbf{q}) - \hat{u}^k(\mathbf{p})) + \epsilon_1 \leq 2\epsilon_0 + \epsilon_1,$$

where the first inequality follows since \mathbf{p} is an ϵ_1 -equilibrium of $\hat{\mathcal{G}}$, hence $\hat{u}^k(\mathbf{p}) - \hat{u}^k(\mathbf{q}) \geq -\epsilon_1$, and the second inequality follows by the lemma's assumption. \square

We turn now to describe a particular class of games that is central in this chapter, the class of potential games (Monderer and Shapley, 1996b).

Definition 2.2.1 (Potential Game). *A potential game is a noncooperative game for which there exists a function $\phi : E \rightarrow \mathbb{R}$ satisfying*

$$\phi(\mathbf{p}^m, \mathbf{p}^{-m}) - \phi(\mathbf{q}^m, \mathbf{p}^{-m}) = u^m(\mathbf{p}^m, \mathbf{p}^{-m}) - u^m(\mathbf{q}^m, \mathbf{p}^{-m}), \quad (2.4)$$

for every $m \in \mathcal{M}$, $\mathbf{p}^m, \mathbf{q}^m \in E^m$, $\mathbf{p}^{-m} \in E^{-m}$. The function ϕ is referred to as a potential function of the game.

Potential games can be regarded as games in which the interests of the players are aligned with a global potential function ϕ . Games that obey condition (2.4) are also known in the literature as *exact* potential games, to distinguish them from other classes of games that relate to a potential function (in a different manner). For simplicity of exposition, we will often write 'potential games' when referring to exact potential games. Potential games have desirable equilibrium properties as summarized in Section 2.5.1.

2.2.2 Games and Flows on Graphs

In noncooperative games, the utility functions capture the preferences of agents at each strategy profile. Note that a Nash equilibrium is defined in terms of payoff differences, suggesting that actual payoffs in the game are not required for the identification of equilibria, as long as the payoff differences are well defined.

A pair of strategy profiles that differ only in the strategy of a single player will be henceforth referred to as *comparable strategy profiles*. We denote the set (of pairs) of comparable strategy profiles by $A \subset E \times E$, i.e., \mathbf{p}, \mathbf{q} are comparable if and only if $(\mathbf{p}, \mathbf{q}) \in A$. A pair of strategy profiles that differ only in the strategy of player m is called a pair of *m -comparable strategy profiles*. The set of pairs of m -comparable strategies is denoted by

$A^m \subset E \times E$. Clearly, $\cup_m A^m = A$, where $A^m \cap A^k = \emptyset$ for any two different players m and k .

For any given m -comparable strategy profiles \mathbf{p} and \mathbf{q} , the payoff difference $[u^m(\mathbf{p}) - u^m(\mathbf{q})]$ will be referred to as their *pairwise comparison*. For any game, we define the *pairwise comparison* function $X : E \times E \rightarrow \mathbb{R}$ as follows

$$X(\mathbf{p}, \mathbf{q}) = \begin{cases} u^m(\mathbf{q}) - u^m(\mathbf{p}) & \text{if } (\mathbf{p}, \mathbf{q}) \text{ are } m\text{-comparable for some } m \in \mathcal{M} \\ 0 & \text{otherwise.} \end{cases} \quad (2.5)$$

In view of Definition 2.2.1, a game is an exact potential game if and only if there exists a function $\phi : E \rightarrow \mathbb{R}$ such that $\phi(\mathbf{q}) - \phi(\mathbf{p}) = X(\mathbf{p}, \mathbf{q})$ for any comparable strategy profiles \mathbf{p} and \mathbf{q} . Note that the pairwise comparisons are uniquely defined for any given game. However, the converse is not true, for instance any two games, for which the utilities differ by a constant have the same pairwise comparisons.

The usual solution concepts in games (e.g., Nash, mixed Nash, correlated equilibria) are defined in terms of pairwise comparisons only. Consequently, games with identical pairwise comparisons share the same equilibrium sets. Thus, in this chapter, we refer to games with identical pairwise comparisons as *strategically equivalent games*².

By employing the notion of pairwise comparisons, we can concisely represent any strategic-form game in terms of a *flow* in a graph. We recall this notion next. Let $G = (N, L)$ be an undirected graph, with set of nodes N and set of links L . An *edge flow* (or just *flow*) on this graph is a function $Y : N \times N \rightarrow \mathbb{R}$ such that $Y(\mathbf{p}, \mathbf{q}) = -Y(\mathbf{q}, \mathbf{p})$ and $Y(\mathbf{p}, \mathbf{q}) = 0$ for $(\mathbf{p}, \mathbf{q}) \notin L$ (Jiang et al., 2011; Bertsimas and Tsitsiklis, 1997). Note that the flow conservation equations are not enforced under this general definition.

Given a game \mathcal{G} , we define a graph where each node corresponds to a pure strategy profile, and each edge connects two comparable strategy profiles. This undirected graph is referred to as the *game graph* and is denoted by $G(\mathcal{G}) \triangleq (E, A)$, where E and A are the strategy profiles and pairs of comparable strategy profiles defined above, respectively. Notice that, by definition, the graph $G(\mathcal{G})$ has the structure of a direct product of M cliques (one per player), with clique m having h_m vertices. The pairwise comparison function $X : E \times E \rightarrow \mathbb{R}$ defines a flow on $G(\mathcal{G})$, as it satisfies $X(\mathbf{p}, \mathbf{q}) = -X(\mathbf{q}, \mathbf{p})$ and $X(\mathbf{p}, \mathbf{q}) = 0$ for $(\mathbf{p}, \mathbf{q}) \notin A$. This flow may thus serve as an equivalent representation of any game (up to a “non-strategic” component). It follows directly from the statements above that two games are strategically equivalent if and only if they have the same flow representation and game graph. An example of the game graph representation is given in Example 2.1.1 of the Section 2.1. Another example can be found below.

²Other strategic equivalence definitions can be found in the literature (see Introduction, for a review of the relevant literature). We focus on this particular definition to ensure that the flow representations of equivalent games are identical, a feature that will be used when studying decompositions of games. It is also true that under our equivalence notion, equivalent games have identical equilibrium sets (mixed, pure and correlated).

Example 2.2.1. *The payoff matrix of the “battle of the sexes” game is given in the following table. The game graph has four vertices, corresponding to the direct product of two 2-cliques, and is presented in Figure 2-2.*




Figure 2-2: Flows on the game graph corresponding to “battle of the sexes”.

The representation of a game as a flow in a graph is natural and useful for the understanding of its strategic interactions, as it abstracts away the absolute utility values and allows for more direct equilibrium-related interpretation. In more mathematical terms, it considers the quotient of the utilities modulo the subspace of games that are “equivalent” to the trivial game (the game where all players receive zero payoff at all strategy profiles), and allows for the identification of “equivalent” games as the same object, a point explored in more detail in later sections. The game graph also contains much structural information. For instance, two games where the directions of arrows are identical (regardless of the flow values) share the same pure Nash equilibria. Our goal in this chapter is to use tools from the theory of graph flows to decompose a game into components, each of which admits tractable equilibrium characterization. The next section provides an overview of the tools that are required for this objective.

2.3 Flows and Helmholtz Decomposition

The objective of this section is to provide a brief overview of the notation and tools required for the analysis of flows on graphs. The basic high-level idea is that under certain conditions (e.g., for graphs arising from games), it is possible to consider graphs as natural topological spaces with nontrivial homological properties. In what follows, we make this idea precise. For simplicity and accessibility to a wider audience, we describe the methods in relatively elementary linear algebraic language, limiting the usage of algebraic topology notions whenever possible. The main technical tool we use is the Helmholtz decomposition theorem, a classical result from algebraic topology with many applications in applied mathematics, including among others electromagnetism, computational geometry and data visualization; see e.g. (Polthier and Preuß, 2002; Tong et al., 2003). In particular, we mention the very interesting recent work by Jiang et al. (Jiang et al., 2011), where the Helmholtz/Hodge decomposition is applied to the problem of statistical ranking for sets of incomplete data.

Consider an undirected graph $G = (E, A)$, where E is the set of the nodes, and A is the

set of edges of the graph³. Since the graph is undirected $(\mathbf{p}, \mathbf{q}) \in A$ if and only if $(\mathbf{q}, \mathbf{p}) \in A$. We denote the set of 3-cliques of the graph by $T = \{(\mathbf{p}, \mathbf{q}, \mathbf{r}) | (\mathbf{p}, \mathbf{q}), (\mathbf{q}, \mathbf{r}), (\mathbf{p}, \mathbf{r}) \in A\}$.

We denote by $C_0 = \{f | f : E \rightarrow \mathbb{R}\}$ the set of real-valued functions on the set of nodes. Recall that the *edge flows* $X : E \times E \rightarrow \mathbb{R}$ are functions which satisfy

$$X(\mathbf{p}, \mathbf{q}) = \begin{cases} -X(\mathbf{q}, \mathbf{p}) & \text{if } (\mathbf{p}, \mathbf{q}) \in A \\ 0 & \text{otherwise.} \end{cases} \quad (2.6)$$

Similarly the *triangular flows* $\Psi : E \times E \times E \rightarrow \mathbb{R}$ are functions for which

$$\Psi(\mathbf{p}, \mathbf{q}, \mathbf{r}) = \Psi(\mathbf{q}, \mathbf{r}, \mathbf{p}) = \Psi(\mathbf{r}, \mathbf{p}, \mathbf{q}) = -\Psi(\mathbf{q}, \mathbf{p}, \mathbf{r}) = -\Psi(\mathbf{p}, \mathbf{r}, \mathbf{q}) = -\Psi(\mathbf{r}, \mathbf{q}, \mathbf{p}), \quad (2.7)$$

and $\Psi(\mathbf{p}, \mathbf{q}, \mathbf{r}) = 0$ if $(\mathbf{p}, \mathbf{q}, \mathbf{r}) \notin T$. Given a graph G , we denote the set of all possible edge flows by C_1 and the set of triangular flows by C_2 . Notice that both C_1 and C_2 are alternating functions of their arguments. It follows from (2.6) that $X(\mathbf{p}, \mathbf{p}) = 0$ for all $X \in C_1$.

The sets C_0 , C_1 and C_2 have a natural structure of vector spaces, with the obvious operations of addition and scalar multiplication. In this chapter, we use the following inner products:

$$\begin{aligned} \langle \phi_1, \phi_2 \rangle_0 &= \sum_{\mathbf{p} \in E} \phi_1(\mathbf{p})\phi_2(\mathbf{p}). \\ \langle X, Y \rangle_1 &= \frac{1}{2} \sum_{(\mathbf{p}, \mathbf{q}) \in A} X(\mathbf{p}, \mathbf{q})Y(\mathbf{p}, \mathbf{q}) \\ \langle \Psi_1, \Psi_2 \rangle_2 &= \sum_{(\mathbf{p}, \mathbf{q}, \mathbf{r}) \in T} \Psi_1(\mathbf{p}, \mathbf{q}, \mathbf{r})\Psi_2(\mathbf{p}, \mathbf{q}, \mathbf{r}). \end{aligned} \quad (2.8)$$

We shall frequently drop the subscript in the inner product notation, as the respective space will often be clear from the context.

We next define linear operators that relate the above defined objects. To that end, let $W : E \times E \rightarrow \mathbb{R}$ be an indicator function for the edges of the graph, namely

$$W(\mathbf{p}, \mathbf{q}) = \begin{cases} 1 & \text{if } (\mathbf{p}, \mathbf{q}) \in A \\ 0 & \text{otherwise.} \end{cases} \quad (2.9)$$

Notice that $W(\mathbf{p}, \mathbf{q})$ can be simply interpreted as the adjacency matrix of the graph G .

The first operator of interest is the *combinatorial gradient operator* $\delta_0 : C_0 \rightarrow C_1$, given by

$$(\delta_0 \phi)(\mathbf{p}, \mathbf{q}) = W(\mathbf{p}, \mathbf{q})(\phi(\mathbf{q}) - \phi(\mathbf{p})), \quad \mathbf{p}, \mathbf{q} \in E, \quad (2.10)$$

for $\phi \in C_0$. An operator which is used in the characterization of ‘‘circulations’’ in edge flows

³The results discussed in this section apply to arbitrary graphs. We use the notation introduced in Section 2.2 since in the rest of the chapter we focus on the game graph introduced there.

is the *curl operator* $\delta_1 : C_1 \rightarrow C_2$, which is defined for all $X \in C_1$ and $\mathbf{p}, \mathbf{q}, \mathbf{r} \in E$ as

$$(\delta_1 X)(\mathbf{p}, \mathbf{q}, \mathbf{r}) = \begin{cases} X(\mathbf{p}, \mathbf{q}) + X(\mathbf{q}, \mathbf{r}) + X(\mathbf{r}, \mathbf{p}) & \text{if } (\mathbf{p}, \mathbf{q}, \mathbf{r}) \in T, \\ 0 & \text{otherwise.} \end{cases} \quad (2.11)$$

We denote the adjoints of the operators δ_0 and δ_1 by δ_0^* and δ_1^* respectively. Recall that given inner products $\langle \cdot, \cdot \rangle_k$ on C_k , the adjoint of δ_k , namely $\delta_k^* : C_{k+1} \rightarrow C_k$, is the unique linear operator satisfying

$$\langle \delta_k f_k, g_{k+1} \rangle_{k+1} = \langle f_k, \delta_k^* g_{k+1} \rangle_k, \quad (2.12)$$

for all $f_k \in C_k$, $g_{k+1} \in C_{k+1}$.

Using the definitions in (2.12), (2.10) and (2.8), it can be readily seen that the adjoint $\delta_0^* : C_1 \rightarrow C_0$ of the combinatorial gradient δ_0 satisfies

$$(\delta_0^* X)(\mathbf{p}) = - \sum_{\mathbf{q} | (\mathbf{p}, \mathbf{q}) \in A} X(\mathbf{p}, \mathbf{q}) = - \sum_{\mathbf{q} \in E} W(\mathbf{p}, \mathbf{q}) X(\mathbf{p}, \mathbf{q}). \quad (2.13)$$

Note that $-(\delta_0^* X)(\mathbf{p})$ represents the total flow “leaving” \mathbf{p} . We shall sometimes refer to the operator $-\delta_0^*$ as the *divergence operator*, due to its similarity to the standard divergence operator in Calculus.

The domains and codomains of the operators $\delta_0, \delta_1, \delta_0^*, \delta_1^*$ are summarized below.

$$\begin{array}{c} C_0 \xrightarrow{\delta_0} C_1 \xrightarrow{\delta_1} C_2 \\ C_0 \xleftarrow{\delta_0^*} C_1 \xleftarrow{\delta_1^*} C_2. \end{array} \quad (2.14)$$

We next define the Laplacian operator, $\Delta_0 : C_0 \rightarrow C_0$, given by

$$\Delta_0 \triangleq \delta_0^* \circ \delta_0, \quad (2.15)$$

where \circ represents operator composition. To simplify the notation, we henceforth omit \circ and write $\Delta_0 = \delta_0^* \delta_0$. Note that functions in C_0 can be represented by vectors of length $|E|$ by indexing all nodes of the graph and constructing a vector whose i th entry is the function evaluated at the i th node. This allows us to easily represent these operators in terms of matrices. In particular, the Laplacian can be expressed as a square matrix of size $|E| \times |E|$; using the definitions for δ_0 and δ_0^* , it follows that

$$[\Delta_0]_{\mathbf{p}, \mathbf{q}} = \begin{cases} \sum_{\mathbf{r} \in E} W(\mathbf{p}, \mathbf{r}) & \text{if } \mathbf{p} = \mathbf{q} \\ -1 & \text{if } \mathbf{p} \neq \mathbf{q} \text{ and } (\mathbf{p}, \mathbf{q}) \in A \\ 0 & \text{otherwise,} \end{cases} \quad (2.16)$$

where, with some abuse of the notation, $[\Delta_0]_{\mathbf{p}, \mathbf{q}}$ denotes the entry of the matrix Δ_0 , with

rows and columns indexed by the nodes \mathbf{p} and \mathbf{q} . The above matrix naturally coincides with the definition of a Laplacian of an undirected graph (Chung, 1997).

Since the entry of $\Delta_0\phi$ corresponding to \mathbf{p} equals $\sum_{\mathbf{q}} W(\mathbf{p}, \mathbf{q})(\phi(\mathbf{p}) - \phi(\mathbf{q}))$, the Laplacian operator gives a measure of the aggregate “value” of a node over all its neighbors. A related operator is

$$\Delta_1 \triangleq \delta_1^* \circ \delta_1 + \delta_0 \circ \delta_0^*, \quad (2.17)$$

known in the literature as the vector Laplacian (Jiang et al., 2011).

We next provide additional flow-related terminology which will be used in association with the above defined operators, and highlight some of their basic properties. In analogy to the well-known identity in vector calculus, $\text{curl} \circ \text{grad} = 0$, we have that δ_0 is a closed form, i.e., $\delta_1 \circ \delta_0 = 0$. An edge flow $X \in C_1$ is said to be *globally consistent* if X corresponds to the combinatorial gradient of some $f \in C_0$, i.e., $X = \delta_0 f$; the function f is referred to as the *potential function* corresponding to X . Equivalently, the set of globally consistent edge flows can be represented as the image of the gradient operator, namely $\text{im}(\delta_0)$. By the closedness of δ_0 , observe that $\delta_1 X = 0$ for every globally consistent edge flow X . We define *locally consistent* edge flows as those satisfying $(\delta_1 X)(\mathbf{p}, \mathbf{q}, \mathbf{r}) = X(\mathbf{p}, \mathbf{q}) + X(\mathbf{q}, \mathbf{r}) + X(\mathbf{r}, \mathbf{p}) = 0$ for all $(\mathbf{p}, \mathbf{q}, \mathbf{r}) \in T$. Note that the kernel of the curl operator $\ker(\delta_1)$ is the set of locally consistent edge flows. The latter subset is generally not equivalent to $\text{im}(\delta_0)$, as there may exist edge flows that are globally inconsistent but locally consistent (in fact, this will happen whenever the graph has a nontrivial topology). We refer to such flows as *harmonic flows*. Note that the operators δ_0, δ_1 are linear operators, thus their image spaces are orthogonal complements of the kernels of their adjoints, i.e., $\text{im}(\delta_0) \perp \ker(\delta_0^*)$ and $\text{im}(\delta_1) \perp \ker(\delta_1^*)$ [similarly, $\text{im}(\delta_0^*) \perp \ker(\delta_0)$ and $\text{im}(\delta_1^*) \perp \ker(\delta_1)$ as can be easily verified using (2.12)].

We state below a basic flow-decomposition theorem, known as the Helmholtz Decomposition⁴, which will be used in our context of noncooperative games. The theorem (see (Jiang et al., 2011)) implies that any graph flow can be decomposed into three orthogonal flows.

Theorem 2.3.1 (Helmholtz Decomposition). *The vector space of edge flows C_1 admits an orthogonal decomposition*

$$C_1 = \text{im}(\delta_0) \oplus \ker(\Delta_1) \oplus \text{im}(\delta_1^*), \quad (2.18)$$

where $\ker(\Delta_1) = \ker(\delta_1) \cap \ker(\delta_0^*)$.

Below we summarize the interpretation of each of the components in the Helmholtz decomposition (see also Figure 2-3):

- $\text{im}(\delta_0)$ – globally consistent flows.

⁴The Helmholtz Decomposition can be generalized to higher dimensions through the Hodge Decomposition theorem (see (Jiang et al., 2011)), however this generalization is not required for our purposes.




Figure 2-3: Helmholtz decomposition of C_1

- $\ker(\Delta_1) = \ker(\delta_1) \cap \ker(\delta_0^*)$ – harmonic flows, which are globally inconsistent but locally consistent. Observe that $\ker(\delta_1)$ consists of locally consistent flows (that may or may not be globally consistent), while $\ker(\delta_0^*)$ consists of globally inconsistent flows (that may or may not be locally consistent).
- $\text{im}(\delta_1^*)$ (or equivalently, the orthogonal complement of $\ker(\delta_1)$) – locally inconsistent flows.

We conclude this section with a brief remark on the decomposition and the flow conservation. For $X \in C_1$, if $\delta_0^* X = 0$, i.e., if for every node, the total flow leaving the node is zero, then we say that X satisfies the flow conservation condition. Theorem 2.3.1 implies that X satisfies this condition only when $X \in \ker(\delta_0^*) = \text{im}(\delta_0)^\perp = \ker(\Delta_1) \oplus \text{im}(\delta_1^*)$. Thus, the flow conservation condition is satisfied for harmonic flows and locally inconsistent flows but not for globally consistent flows.

2.4 Canonical Decomposition of Games

In this section we obtain a canonical decomposition of an arbitrary game into basic components, by combining the game graph representation introduced in Section 2.2.2 with the Helmholtz decomposition discussed above.

Section 2.4.1 introduces the relevant operators that are required for formulating the results. In Section 2.4.2 we provide the basic decomposition theorem, which states that the space of games can be decomposed as a direct sum of three subspaces, referred to as the *potential*, *harmonic* and *nonstrategic* subspaces. In Section 2.4.3, we focus on bimatrix games, and provide explicit expressions for the decomposition.

2.4.1 Preliminaries

We consider a game \mathcal{G} with set of players \mathcal{M} , strategy profiles $E \triangleq E^1 \times \dots \times E^M$, and game graph $G(\mathcal{G}) = (E, A)$. Using the notation of the previous section, the utility functions of each player can be viewed as elements of C_0 , i.e., $u^m \in C_0$ for all $m \in \mathcal{M}$. For given \mathcal{M} and E , every game is uniquely defined by its set of utility functions. Hence, the space of games with players \mathcal{M} and strategy profiles E can be identified as $\mathcal{G}_{\mathcal{M}, E} \cong C_0^M$. In the rest of the chapter we use the notations $\{u^m\}_{m \in \mathcal{M}}$ and $\mathcal{G} = \langle \mathcal{M}, \{E^m\}, \{u^m\} \rangle$ interchangeably when referring to games.

The pairwise comparison function $X(\mathbf{p}, \mathbf{q})$ of a game, defined in (2.5), corresponds to a flow on the game graph, and hence it belongs to C_1 . In general, the flows representing games have some special structure. For example, the pairwise comparison between any two comparable strategy profiles is associated with the payoff of exactly a single player. It is therefore required to introduce *player-specific* operators and highlight some important identities between them, as we elaborate below.

Let $W^m : E \times E \rightarrow \mathbb{R}$ be the indicator function for m -comparable strategy profiles, namely

$$W^m(\mathbf{p}, \mathbf{q}) = \begin{cases} 1 & \text{if } \mathbf{p}, \mathbf{q} \text{ are } m\text{-comparable} \\ 0 & \text{otherwise.} \end{cases}$$

Recalling that any pair of strategy profiles cannot be comparable by more than a single player, we have

$$W^m(\mathbf{p}, \mathbf{q})W^k(\mathbf{p}, \mathbf{q}) = 0, \quad \text{for all } k \neq m \text{ and } \mathbf{p}, \mathbf{q} \in E, \quad (2.19)$$

and

$$W = \sum_{m \in \mathcal{M}} W^m, \quad (2.20)$$

where W is the indicator function of comparable strategy profiles (edges of the game graph) defined in (2.9). Note that this can be interpreted as a decomposition of the adjacency matrix of $G(\mathcal{G})$, where the different components correspond to the edges associated with different players.

Given $\phi \in C_0$, we define $D_m : C_0 \rightarrow C_1$ such that

$$(D_m \phi)(\mathbf{p}, \mathbf{q}) = W^m(\mathbf{p}, \mathbf{q}) (\phi(\mathbf{q}) - \phi(\mathbf{p})). \quad (2.21)$$

This operator quantifies the change in ϕ between strategy profiles that are m -comparable. Using this operator, we can represent the pairwise differences X of a game with payoffs $\{u^m\}_{m \in \mathcal{M}}$ as follows:

$$X = \sum_{m \in \mathcal{M}} D_m u^m. \quad (2.22)$$

We define a relevant operator $D : C_0^M \rightarrow C_1$, such that $D = [D_1 \dots, D_M]$. As can be seen from (2.22), for a game with collection of utilities $u = [u^1; u^2 \dots; u^M] \in C_0^M$, the pairwise differences can alternatively be represented by Du .

Let $\Lambda_m : C_1 \rightarrow C_1$ be an operator so that

$$(\Lambda_m X)(\mathbf{p}, \mathbf{q}) = W^m(\mathbf{p}, \mathbf{q})X(\mathbf{p}, \mathbf{q})$$

for every $X \in C_1$, $\mathbf{p}, \mathbf{q} \in E$. From (2.20), it can be seen that for any $X \in C_1$, $\sum_{m \in \mathcal{M}} \Lambda_m X = X$. The definition of Λ_m and (2.19) imply that $\Lambda_m \Lambda_k = 0$ for $k \neq m$. Additionally, the definition of the inner product in C_1 implies that for $X, Y \in C_1$, it follows that $\langle \Lambda_m X, Y \rangle = \langle X, \Lambda_m Y \rangle$, i.e., Λ_m is self-adjoint.

This operator provides a convenient description for the operator D_m . From the definitions of D_m and Λ_m , it immediately follows that $D_m = \Lambda_m \delta_0$, and since $\sum_{m \in \mathcal{M}} \Lambda_m X = X$ for all $X \in C_1$,

$$\delta_0 = \sum_{m \in \mathcal{M}} \Lambda_m \delta_0 = \sum_{m \in \mathcal{M}} D_m.$$

Since Λ_m is self-adjoint, the adjoint of D_m , which is denoted by $D_m^* : C_1 \rightarrow C_0$, is given by:

$$D_m^* = \delta_0^* \Lambda_m.$$

Using (2.13) and the above definitions, it follows that

$$(D_m^* X)(\mathbf{p}) = - \sum_{\mathbf{q} \in E} W^m(\mathbf{p}, \mathbf{q})X(\mathbf{p}, \mathbf{q}), \quad \text{for all } X \in C_1, \quad (2.23)$$

and

$$\delta_0^* = \sum_{m \in \mathcal{M}} D_m^*. \quad (2.24)$$

Observe that $D_k^* D_m = \delta_0^* \Lambda_k \Lambda_m \delta_0 = 0$ for $k \neq m$. This immediately implies that the image spaces of $\{D_m\}_{m \in \mathcal{M}}$ are orthogonal, i.e., $D_k^* D_m = 0$. Let D_m^\dagger denote the (Moore-Penrose) pseudoinverse of D_m , with respect to the inner products introduced in Section 2.3. By the properties of the pseudoinverse, we have $\ker D_m^\dagger = (\text{im } D_m)^\perp$. Thus, orthogonality of the image spaces of D_k operators imply that $D_k^\dagger D_m = 0$ for $k \neq m$.

The orthogonality leads to the following expression for the Laplacian operator,

$$\Delta_0 = \delta_0^* \delta_0 = \sum_{k \in \mathcal{M}} D_k^* \sum_{m \in \mathcal{M}} D_m = \sum_{m \in \mathcal{M}} D_m^* D_m. \quad (2.25)$$

In view of (2.21) and (2.23), D_m and $-D_m^*$ are the gradient and divergence operators on the graph of m -comparable strategy profiles (E, A^m) . Therefore, the operator $\Delta_{0,m} \triangleq D_m^* D_m$ is the Laplacian of the graph induced by m -comparable strategies, and is referred to as the *Laplacian operator of the m -comparable strategy profiles*. It follows from (2.25)

that

$$\Delta_0 = \sum_{m \in \mathcal{M}} \Delta_{0,m}.$$

The relation between the Laplacian operators Δ_0 and $\Delta_{0,m}$ is illustrated in Figure 2-4.




Figure 2-4: A game with two players, each of which has three strategies. A node (i, j) represents a strategy profile in which player 1 and player 2 use strategies i and j , respectively. The Laplacian $\Delta_{0,1}$ ($\Delta_{0,2}$) is defined on the graph whose edges are represented by dashed (solid) lines. The Laplacian Δ_0 is defined on the graph that includes all edges.

Similarly, $\delta_1 \Lambda_m$ is the curl operator associated with the subgraph (E, A^m) . From the closedness of the curl ($\delta_1 \Lambda_m$) and gradient ($\Lambda_m \delta_0$) operators defined on this subgraph, we obtain $\delta_1 \Lambda_m^2 \delta_0 = 0$. Observing that $\Lambda_m^2 \delta_0 = \Lambda_m \delta_0 = D_m$, it follows that

$$\delta_1 D_m = 0. \quad (2.26)$$

This result also implies that $\delta_1 D = 0$, i.e., the pairwise comparisons of games belong to $\ker \delta_1$. Thus, it follows from Theorem 2.3.1 that the pairwise comparisons do not have a locally inconsistent component. Intuitively, there is no local inconsistency, because only three cliques in the game graph are due to unilateral deviations of a single player, and hence cannot lead to local inconsistency.

Lastly, we introduce projection operators that will be useful in the subsequent analysis. Consider the operator,

$$\Pi_m = D_m^\dagger D_m.$$

Since D_m is a linear operator, Π_m is a projection operator⁵ to the orthogonal complement of the kernel of D_m . Using these operators, we define $\Pi : C_0^M \rightarrow C_0^M$ such that $\Pi = \text{diag}(\Pi_1, \dots, \Pi_M)$, i.e., for $u = \{u_m\}_{m \in \mathcal{M}} \in C_0^M$, we have $\Pi u = [\Pi_1 u^1; \dots \Pi_M u^M] \in C_0^M$. We extend the inner product in C_0 to C_0^M (by defining the inner product as the sum of the

⁵For any linear operator L , $L^\dagger L$ is a projection operator on the orthogonal complement of the kernel of L (see (Golub and Van Loan, 1996)).

inner products in all C_0 components), and denote by D^\dagger the pseudoinverse of D according to this inner product. In Lemma 2.4.4, we will show that Π is equivalent to the projection operator to the orthogonal complement of the kernel of D , i.e., $\Pi = D^\dagger D$.

For easy reference, Table 2.1 provides a summary of notation. We next state some basic facts about the operators we introduced, which will be used in the subsequent analysis. The proofs of these results can be found in Section 2.8.

\mathcal{G}	A game instance $\langle \mathcal{M}, \{E^m\}_{m \in \mathcal{M}}, \{u^m\}_{m \in \mathcal{M}} \rangle$.
\mathcal{M}	Set of players, $\{1, \dots, M\}$.
E^m	Set of actions for player m , $E^m = \{1, \dots, h_m\}$.
E	Joint action space $\prod_{m \in \mathcal{M}} E^m$.
u^m	Utility function of player m . We have $u^m \in C_0$.
W^m	Indicator function for m -comparable strategy profiles, $W^m : E \times E \rightarrow \{0, 1\}$.
W	A function indicating whether strategy profiles are comparable, $W : E \times E \rightarrow \{0, 1\}$.
C_0	Space of utilities, $C_0 = \{u^m u^m : E \rightarrow \mathbb{R}\}$. Note that $C_0 \cong \mathbb{R}^{ E }$.
C_1	Space of pairwise comparison functions from $E \times E$ to \mathbb{R} .
δ_0	Gradient operator, $\delta_0 : C_0 \rightarrow C_1$, satisfying $(\delta_0 \phi)(\mathbf{p}, \mathbf{q}) = W(\mathbf{p}, \mathbf{q}) (\phi(\mathbf{q}) - \phi(\mathbf{p}))$.
D_m	$D_m : C_0 \rightarrow C_1$, such that $(D_m \phi)(\mathbf{p}, \mathbf{q}) = W^m(\mathbf{p}, \mathbf{q}) (\phi(\mathbf{q}) - \phi(\mathbf{p}))$.
D	$D : C_0^M \rightarrow C_1$, such that $D(u^1; \dots; u^M) = \sum_{m \in \mathcal{M}} D_m u^m$.
δ_0^*, D_m^*	$\delta_0^*, D_m^* : C_1 \rightarrow C_0$ are the adjoints of the operators δ_0 and D_m , respectively.
Δ_0	Laplacian for the game graph. $\Delta_0 : C_0 \rightarrow C_0$; satisfies $\Delta_0 = \delta_0^* \delta_0 = \sum_{m \in \mathcal{M}} \Delta_{0,m}$.
$\Delta_{0,m}$	Laplacian for the graph of m -comparable strategies, $\Delta_{0,m} : C_0 \rightarrow C_0$; satisfies $\Delta_{0,m} = D_m^* D_m = D_m^* \delta_0$.
Π_m	Projection operator onto the orthogonal complement of kernel of D_m , $\Pi_m : C_0 \rightarrow C_0$; satisfies $\Pi_m = D_m^\dagger D_m$.

Table 2.1: Notation summary

Lemma 2.4.1. *The Laplacian of the graph induced by m -comparable strategies and the projection operator Π_m are related by $\Delta_{0,m} = h_m \Pi_m$, where $h_m = |E^m|$ denotes the number of strategies of player m .*

Lemma 2.4.2. *The kernels of operators D_m , Π_m and $\Delta_{0,m}$ coincide, namely $\ker(D_m) = \ker(\Pi_m) = \ker(\Delta_{0,m})$. Furthermore, a basis for these kernels is given by a collection $\{\nu_{\mathbf{q}^{-m}}\}_{\mathbf{q}^{-m} \in E^{-m}} \in C_0$ such that*

$$\nu_{\mathbf{q}^{-m}}(\mathbf{p}) = \begin{cases} 1 & \text{if } \mathbf{p}^{-m} = \mathbf{q}^{-m} \\ 0 & \text{otherwise} \end{cases} \quad (2.27)$$

Lemma 2.4.3. *The Laplacian Δ_0 of the game graph (the graph of comparable strategy profiles) always has eigenvalue 0 with multiplicity 1, corresponding to the constant eigenfunction (i.e., $f \in C_0$ such that $f(\mathbf{p}) = 1$ for all $\mathbf{p} \in E$).*

Lemma 2.4.4. *The pseudoinverses of operators D_m and D satisfy the following identities:*

1. $D_m^\dagger = \frac{1}{h_m} D_m^*$,
2. $(\sum_{i \in \mathcal{M}} D_i)^\dagger D_j = (\sum_{i \in \mathcal{M}} D_i^* D_i)^\dagger D_j^* D_j$,
3. $D^\dagger = [D_1^\dagger; \dots; D_M^\dagger]$,
4. $\Pi = D^\dagger D$,
5. $DD^\dagger \delta_0 = \delta_0$.

2.4.2 Decomposition of Games

In this subsection we prove that the space of games $\mathcal{G}_{\mathcal{M}, E}$ is a direct sum of three subspaces – potential, harmonic and nonstrategic, each with distinguishing properties.

We start our discussion by formalizing the notion of nonstrategic information. Consider two games $\mathcal{G}, \hat{\mathcal{G}} \in \mathcal{G}_{\mathcal{M}, E}$ with utilities $\{u^m\}_{m \in \mathcal{M}}$ and $\{\hat{u}^m\}_{m \in \mathcal{M}}$ respectively. Assume that the utility functions $\{u^m\}_{m \in \mathcal{M}}$ satisfy $u^m(\mathbf{p}^m, \mathbf{p}^{-m}) = \hat{u}^m(\mathbf{p}^m, \mathbf{p}^{-m}) + \alpha(\mathbf{p}^{-m})$ where α is an arbitrary function. It can be seen that these two games have exactly the same pairwise comparison functions, hence they are strategically equivalent. To fix a representative for strategically equivalent games, we introduce below a notion of normalization for games.

Definition 2.4.1 (Normalized games). *We say that a game with utility functions $\{u^m\}_{m \in \mathcal{M}}$ is normalized or does not contain nonstrategic information if*

$$\sum_{\mathbf{p}^m \in E^m} u^m(\mathbf{p}^m, \mathbf{p}^{-m}) = 0 \quad (2.28)$$

for all $\mathbf{p}^{-m} \in E^{-m}$ and all $m \in \mathcal{M}$.

It will be shown in the sequel that for each game there is a unique strategically equivalent normalized game, hence normalized games can be used to identify representatives for strategically equivalent games. Normalization can be made with an arbitrary constant. However, in order to simplify the subsequent analysis we normalize the sum of the payoffs to zero. Intuitively, in normalized games, given the strategies of a player's opponents, the expected payoff of a player for a uniformly mixed strategy is equal to zero. The following lemma characterizes the set of normalized games in terms of the operators introduced in the previous section.

Lemma 2.4.5. *Given a game \mathcal{G} with utilities $u = \{u^m\}_{m \in \mathcal{M}}$, the following are equivalent: (i) \mathcal{G} is normalized, (ii) $\Pi_m u^m = u^m$ for all m , (iii) $\Pi u = u$, (iv) $u \in (\ker D)^\perp$.*

Proof. The equivalence of (iii) and (iv) is immediate since by Lemma 2.4.4, $\Pi = D^\dagger D$ is a projection operator to the orthogonal complement of the kernel of D . The equivalence of (ii) and (iii) follows from the definition of $\Pi = \text{diag}(\Pi_1, \dots, \Pi_M)$. To complete the proof we prove (i) and (ii) are equivalent.

Observe that (2.28) holds if and only if $\langle u^m, \nu_{\mathbf{q}^{-m}}(\mathbf{p}) \rangle = 0$ for all $\mathbf{q}^{-m} \in E^{-m}$, where $\nu_{\mathbf{q}^{-m}}$ is as defined in (2.27). Lemma 2.4.2 implies that $\{\nu_{\mathbf{q}^{-m}}\}$ are basis vectors of $\ker D_m$. Thus, it follows that (2.28) holds if and only if u^m is orthogonal to all of the basis vectors of $\ker D_m$, or equivalently when $u^m \in (\ker D_m)^\perp$. Since $\Pi_m = D_m^\dagger D_m$ is a projection operator to $(\ker D_m)^\perp$, we have $u^m \in (\ker D_m)^\perp$ if and only if $\Pi_m u^m = u^m$, and the claim follows. \square

Using Lemma 2.4.5, we next show below that for each game \mathcal{G} there exists a unique strategically equivalent game which is normalized (contains no nonstrategic information).

Lemma 2.4.6. *Let \mathcal{G} be a game with utilities $\{u^m\}_{m \in \mathcal{M}}$. Then there exists a unique game $\hat{\mathcal{G}}$ which (i) has the same pairwise comparison function as \mathcal{G} and (ii) is normalized. Moreover the utilities $\hat{u} = \{\hat{u}^m\}_{m \in \mathcal{M}}$ of $\hat{\mathcal{G}}$ satisfy $\hat{u}^m = \Pi_m u^m$ for all m .*

Proof. To prove the claim we show that given $u = \{u^m\}_{m \in \mathcal{M}}$, the game with the collection of utilities $D^\dagger D u = \Pi u$, is a normalized game with the same pairwise comparisons, and moreover there cannot be another normalized game which has the same pairwise comparisons.

Since Π is a projection operator, it follows that $\Pi \Pi u = \Pi u$, and hence, Lemma 2.4.5 implies that Πu is normalized. Additionally, using properties of the pseudoinverse we have $D \Pi u = D D^\dagger D u = D u$, thus Πu and u have the same pairwise comparison.

Let $v \in C_0^M$ denote the collection of payoff functions of a game which is normalized and has the same pairwise comparison as u . It follows that $D v = D u = D \Pi u$, and hence $v - \Pi u \in \ker D$. On the other hand, since both v and Πu are normalized, by Lemma 2.4.5, we have $v, \Pi u \in (\ker D)^\perp$, and thus $v - \Pi u \in (\ker D)^\perp$. Therefore, it follows that $v - \Pi u = 0$, hence Πu is the collection of utility functions of the unique normalized game, which has the same pairwise comparison function as \mathcal{G} . By Lemma 2.4.4, $\Pi u = \{\Pi_m u^m\}$, hence the claim follows. \square

We are now ready to define the subspaces of games that will appear in our decomposition result.

Definition 2.4.2. *The potential subspace \mathcal{P} , the harmonic subspace \mathcal{H} and the nonstrategic subspace \mathcal{N} are defined as:*

$$\begin{aligned} \mathcal{P} &\triangleq \{u \in C_0^M \mid u = \Pi u \text{ and } D u \in \text{im } \delta_0\} \\ \mathcal{H} &\triangleq \{u \in C_0^M \mid u = \Pi u \text{ and } D u \in \ker \delta_0^*\} \\ \mathcal{N} &\triangleq \{u \in C_0^M \mid u \in \ker D\}. \end{aligned} \tag{2.29}$$

Since the operators involved in the above definitions are linear, it follows that the sets \mathcal{P} , \mathcal{H} and \mathcal{N} are indeed subspaces.

Lemma 2.4.5 implies that the games in \mathcal{P} and \mathcal{H} are normalized (contain no nonstrategic information). The flows generated by the games in these two subspaces are related to the

flows induced by the Helmholtz decomposition. It follows from the definitions that the flows generated by a game in \mathcal{P} are in the image space of δ_0 and the flows generated by a game in \mathcal{H} are in the kernel of δ_0^* . Thus, \mathcal{P} corresponds to the set of normalized games, which have globally consistent pairwise comparisons. Due to (2.26), the pairwise comparisons of games do not have locally inconsistent components, thus Theorem 2.3.1 implies that \mathcal{H} corresponds to the set of normalized games, which have globally inconsistent but locally consistent pairwise comparisons. Hence, from the perspective of the Helmholtz decomposition, the flows generated by games in \mathcal{P} and \mathcal{H} are gradient and harmonic flows respectively. On the other hand the flows generated by games in \mathcal{N} are always zero, since $Du = 0$ in such games.

As discussed in the previous section the image spaces of D_m are orthogonal. Thus, since by definition $Du = \sum_{m \in \mathcal{M}} D_m u^m$, it follows that $u = \{u^m\}_{m \in \mathcal{M}} \in \ker D$ if and only if $u^m \in \ker D_m$ for all $m \in \mathcal{M}$. Using these facts together with Lemma 2.4.5, we obtain the following alternative description of the subspaces of games:

$$\begin{aligned} \mathcal{P} &= \{ \{u^m\}_{m \in \mathcal{M}} \mid D_m u^m = D_m \phi \text{ and } \Pi_m u^m = u^m \text{ for all } m \in \mathcal{M} \text{ and some } \phi \in C_0 \} \\ \mathcal{H} &= \{ \{u^m\}_{m \in \mathcal{M}} \mid \delta_0^* \sum_{m \in \mathcal{M}} D_m u^m = 0 \text{ and } \Pi_m u^m = u^m \text{ for all } m \in \mathcal{M} \} \\ \mathcal{N} &= \{ \{u^m\}_{m \in \mathcal{M}} \mid D_m u^m = 0 \text{ for all } m \in \mathcal{M} \}. \end{aligned} \tag{2.30}$$

The main result of this section shows that not only these subspaces have distinct properties in terms of the flows they generate, but in fact they form a direct sum decomposition of the space of games. We exploit the Helmholtz decomposition (Theorem 2.3.1) for the proof.

Theorem 2.4.1. *The space of games $\mathcal{G}_{\mathcal{M},E}$ is a direct sum of the potential, harmonic and nonstrategic subspaces, i.e., $\mathcal{G}_{\mathcal{M},E} = \mathcal{P} \oplus \mathcal{H} \oplus \mathcal{N}$. In particular, given a game with utilities $u = \{u^m\}_{m \in \mathcal{M}}$, it can be uniquely decomposed in three components:*

- **Potential Component:** $u_P \triangleq D^\dagger \delta_0 \delta_0^\dagger Du$
- **Harmonic Component:** $u_H \triangleq D^\dagger (I - \delta_0 \delta_0^\dagger) Du$
- **Nonstrategic Component:** $u_N \triangleq (I - D^\dagger D)u$

where $u_P + u_H + u_N = u$, and $u_P \in \mathcal{P}$, $u_H \in \mathcal{H}$, $u_N \in \mathcal{N}$. The potential function associated with u_P is $\phi \triangleq \delta_0^\dagger Du$.

Proof. The decomposition of $\mathcal{G}_{\mathcal{M},E}$ described above follows directly from pulling back the Helmholtz decomposition of C_1 through the map D , and removing the kernel of D ; see Figure 2-5.

The components of the decomposition clearly satisfy $u_P + u_H + u_N = u$. We verify the inclusion properties, according to (2.29). Both u_P and u_H are orthogonal to $\mathcal{N} = \ker D$, since they are in the range of D^\dagger .

$$\begin{array}{ccccc}
\mathcal{G}_{\mathcal{M},E} \cong C_0^M & & & & \\
& \searrow D & & & \\
C_0 & \xrightarrow{\delta_0} & C_1 & \xrightarrow{\delta_1} & C_2
\end{array}$$

Figure 2-5: The Helmholtz decomposition of the space of flows (C_1) can be pulled back through D to a direct sum decomposition of the space of games ($\mathcal{G}_{\mathcal{M},E}$).

- For the potential component, let $\phi \in C_0$ be such that $\phi = \delta_0^\dagger Du$. Then, we have $Du_P \in \text{im}(\delta_0)$, since

$$Du_P = DD^\dagger \delta_0 \delta_0^\dagger Du = \delta_0 \delta_0^\dagger Du = \delta_0 \phi,$$

where we used the definition of u_P , the property (v) in Lemma 2.4.4 and the definition of ϕ , respectively. This equality also implies that ϕ is the potential function associated with u_P .

- For the harmonic component u_H , we have $Du_H \in \ker \delta_0^*$:

$$\delta_0^* Du_H = \delta_0^* DD^\dagger (I - \delta_0 \delta_0^\dagger) Du = \delta_0^* (I - \delta_0 \delta_0^\dagger) Du = 0,$$

as follows from the definition of u_H , the property (v) in Lemma 2.4.4, and properties of the pseudoinverse.

- To check that $u_N \in \mathcal{N}$, we have

$$Du_N = D(I - D^\dagger D)u = (D - DD^\dagger D)u = 0.$$

In order to prove that the direct sum decomposition property holds, we assume that there exists $\hat{u}_P \in \mathcal{P}$, $\hat{u}_H \in \mathcal{H}$ and $\hat{u}_N \in \mathcal{N}$ such that $\hat{u}_P + \hat{u}_H + \hat{u}_N = 0$. Observe that $I - D^\dagger D$ is a projection operator to the kernel of D . Thus, from the definition of the subspaces \mathcal{P} , \mathcal{H} and \mathcal{N} , it follows that $(I - D^\dagger D)\hat{u}_N = \hat{u}_N$ and $(I - D^\dagger D)\hat{u}_P = (I - D^\dagger D)\hat{u}_H = 0$. Similarly, $\delta_0 \delta_0^\dagger$ is a projection operator to the image of δ_0 . Since by definition $D\hat{u}_P \in \text{im} \delta_0$, and $D\hat{u}_H \in \ker \delta_0^* = (\text{im} \delta_0)^\perp$, it follows that $\delta_0 \delta_0^\dagger D\hat{u}_P = D\hat{u}_P$ and $\delta_0 \delta_0^\dagger D\hat{u}_H = 0$.

Using these identities, it follows that

$$\begin{aligned}
(D^\dagger \delta_0 \delta_0^\dagger D)(\hat{u}_P + \hat{u}_H + \hat{u}_N) &= \hat{u}_P \\
D^\dagger (I - \delta_0 \delta_0^\dagger) D(\hat{u}_P + \hat{u}_H + \hat{u}_N) &= \hat{u}_H \\
(I - D^\dagger D)(\hat{u}_P + \hat{u}_H + \hat{u}_N) &= \hat{u}_N,
\end{aligned}$$

Since, $\hat{u}_P + \hat{u}_H + \hat{u}_N = 0$ by our assumption, it follows that $\hat{u}_P = \hat{u}_H = \hat{u}_N = 0$, and hence the direct sum decomposition property follows. \square

The pseudoinverse of a linear operator L , projects its argument to the image space of L , and then, pulls the projection back to the the domain of L . Thus, intuitively, the potential function $\phi = \delta_0^\dagger Du$, defined in the theorem, is such that the gradient flow associated with it ($\delta_0 \phi$) approximates the flow in the original game (Du), in the best possible way. The potential component of the game can be identified by pulling back this gradient flow through D to C_0^M . The harmonic component can similarly be obtained using the harmonic flow.

Since $\delta_0 = \sum_{m \in \mathcal{M}} D_m$, it follows that $\phi = \delta_0^\dagger Du = (\sum_{m \in \mathcal{M}} D_m)^\dagger \sum_{m \in \mathcal{M}} D_m u^m$. Thus, Lemma 2.4.4 (ii), and identities $\Delta_{0,m} = D_m^* D_m$ and $\Delta_0 = \sum_{m \in \mathcal{M}} \Delta_{0,m}$ imply that

$$\phi = \Delta_0^\dagger \sum_{m \in \mathcal{M}} \Delta_{0,m} u^m.$$

Additionally, from Lemma 2.4.4 (iii) and (iv) it follows that $D^\dagger \delta_0 = [D_1^\dagger D_1; \dots; D_M^\dagger D_M] = [\Pi_1; \dots; \Pi_M]$ and $D^\dagger D = \Pi = \text{diag}(\Pi_1, \dots, \Pi_M)$. Using these identities, the utility functions of components of a game can alternatively be expressed as follows:

- **Potential Component:** $u_P^m = \Pi_m \phi$, for all $m \in \mathcal{M}$,
- **Harmonic Component:** $u_H^m = \Pi_m u^m - \Pi_m \phi$, for all $m \in \mathcal{M}$,
- **Nonstrategic Component:** $u_N^m = (I - \Pi_m) u^m$, for all $m \in \mathcal{M}$.

It can be seen that the definitions of the subspaces do not rely on the inner product in C_0^M . Thus, the direct sum property implies that the decomposition is canonical, i.e., it is independent of the inner product used in C_0^M . The above expressions provide closed form solutions for the utility functions in the decomposition, without reference to this inner product. We show in Section 2.6 that our decomposition is indeed orthogonal with respect to a natural inner product in C_0^M .

Note that $\Delta_0 : C_0 \rightarrow C_0$, whereas $\delta_0 : C_0 \rightarrow C_1$. Since C_1 and C_0 are associated with the edges and the nodes of the game graph respectively, in general C_1 is higher dimensional than C_0 . Therefore, calculating Δ_0^\dagger is computationally more tractable than calculating δ_0^\dagger . Hence, the alternative expressions for the components of a game and the potential function ϕ , have computational benefits over using the results of Theorem 2.4.1 directly.

In (Facchini et al., 1997) and (Voorneveld et al., 1999), decompositions of potential games to a ‘‘congestion’’ component (where all players have identical utility functions), and a ‘‘dummy’’ component (which is nonstrategic) were provided. We note that this decomposition is different from the decomposition suggested by Theorem 2.4.1 even when the original game is a potential game. In particular, in the potential component of a game, players can have different utility functions (as can be seen from the alternative expressions,

for utilities in different components of a game). This is not the case for the “congestion component” provided in these works.

We conclude this section by characterizing the dimensions of the potential, harmonic and nonstrategic subspaces. The proof can be found in the Section 2.8.

Proposition 2.4.1. *The dimensions of the subspaces \mathcal{P} , \mathcal{H} and \mathcal{N} are:*

1. $\dim(\mathcal{P}) = \prod_{m \in \mathcal{M}} h_m - 1$,
2. $\dim(\mathcal{H}) = (M - 1) \prod_{m \in \mathcal{M}} h_m - \sum_{m \in \mathcal{M}} \prod_{k \neq m} h_k + 1$.
3. $\dim(\mathcal{N}) = \sum_{m \in \mathcal{M}} \prod_{k \neq m} h_k$.

2.4.3 An Example: Decomposition of Bimatrix Games

We conclude this section by providing an explicit decomposition result for bimatrix games, i.e., finite games with two players. Consider a bimatrix game, where the payoff matrix of the row player is given by A , and that of the column player is given by B ; that is, when the row player plays i and the column player plays j , the row player’s payoff is equal to A_{ij} and the column player’s payoff is equal to B_{ij} .

Assume that both the row player and the column player have the same number h of strategies. It immediately follows from Proposition 2.4.1 that $\dim \mathcal{P} = h^2 - 1$, $\dim \mathcal{H} = (h - 1)^2$ and $\dim \mathcal{N} = 2h$. For simplicity, we further assume that the payoffs are normalized⁶. Thus, the definition of normalized games implies that $\mathbf{1}^T A = B \mathbf{1} = 0$, where $\mathbf{1}$ denotes the vector of ones. Denote by A_P (B_P) and A_H (B_H) respectively, the payoff matrices of the row player (column player) in the potential and harmonic components of the game. Using our decomposition result (Theorem 2.4.1), it follows that

$$(A_P, B_P) = (S + \Gamma, S - \Gamma), \quad (A_H, B_H) = (D - \Gamma, -D + \Gamma), \quad (2.31)$$

where $S = \frac{1}{2}(A + B)$, $D = \frac{1}{2}(A - B)$, $\Gamma = \frac{1}{2h}(A \mathbf{1} \mathbf{1}^T - \mathbf{1} \mathbf{1}^T B)$. Interestingly, the payoff of each player in these components depends on the payoffs of both players in the original game: The potential component of the game relates to the average of the payoffs in the original game and the harmonic component relates to the difference in payoffs of players. The Γ term ensures that the potential and harmonic components do not contain nonstrategic information. We use the above characterization in the next example for obtaining explicit payoff matrices for each of the game components.

Example 2.4.1 (Generalized Rock-Paper-Scissors). *The payoff matrix of the generalized Rock-Paper-Scissors (RPS) game is given in Table 2.2a. Tables 2.2b, 2.2c and 2.2d include the nonstrategic, potential and the harmonic components of the game. The special case*

⁶Lemma 2.4.2 and Lemma 2.4.6 imply that if the payoffs are not normalized, the normalized payoffs can be obtained as $(A - \frac{1}{h} \mathbf{1} \mathbf{1}^T A, B - \frac{1}{h} B \mathbf{1} \mathbf{1}^T)$.

where $x = y = z = \frac{1}{3}$ corresponds to the celebrated RPS game. Note that in this case, the potential component of the game is equal to zero.

	R	P	S
R	0, 0	-3x, 3x	3y, -3y
P	3x, -3x	0, 0	-3z, 3z
S	-3y, 3y	3z, -3z	0, 0

(a) Generalized RPS Game

	R	P	S
R	(x - y), (x - y)	(z - x), (x - y)	(y - z), (x - y)
P	(x - y), (z - x)	(z - x), (z - x)	(y - z), (z - x)
S	(x - y), (y - z)	(z - x), (y - z)	(y - z), (y - z)

(b) Nonstrategic Component

	R	P	S
R	(y - x), (y - x)	(y - x), (x - z)	(y - x), (z - y)
P	(x - z), (y - x)	(x - z), (x - z)	(x - z), (z - y)
S	(z - y), (y - x)	(z - y), (x - z)	(z - y), (z - y)

(c) Potential Component

	R	P	S
R	0, 0	-(x + y + z), (x + y + z)	(x + y + z), -(x + y + z)
P	(x + y + z), -(x + y + z)	0, 0	-(x + y + z), (x + y + z)
S	-(x + y + z), (x + y + z)	(x + y + z), -(x + y + z)	0, 0

(d) Harmonic Component

Table 2.2: Generalized RPS game and its components.

2.5 Potential and Harmonic Games

In this section we study the classes of games that are naturally motivated by our decomposition. In particular, we focus on two classes of games: (i) Games with no harmonic component, (ii) Games with no potential component. We show that the first class is equivalent to the well-known class of *potential games*. We refer to the games in the second class as *harmonic games*. Pictorially, we have

$$\underbrace{\mathcal{P}}_{\text{Potential games}} \oplus \underbrace{\mathcal{N} \oplus \mathcal{H}}_{\text{Harmonic games}}.$$

In Sections 2.5.1 and 2.5.2, we establish the equivalence of the subspace definitions of potential and harmonic games, and the utility definitions given in equations (2.1) and (2.2). Additionally we develop and discuss several properties of these classes of games, with particular emphasis on their equilibria. Since potential games have been extensively studied in the literature, our main focus is on harmonic games. The nonstrategic component does not have an impact on the strategic actions of players, however, we establish in 2.5.3

	Potential Games	Harmonic Games
Subspaces	$\mathcal{P} \oplus \mathcal{N}$	$\mathcal{H} \oplus \mathcal{N}$
Flows	Globally consistent	Locally consistent but globally inconsistent
Pure NE	Always Exists	Generically does not exist
Mixed NE	Always Exists	-Uniformly mixed strategy is always a mixed NE -Players do not strictly prefer their equilibrium strategies.
Special Cases	-	-(two players) Set of mixed Nash equilibria coincides with the set of correlated equilibria -(two players & equal number of strategies) Uniformly mixed strategy is the unique mixed NE

Table 2.3: Properties of potential and harmonic games.

that the nonstrategic component affects efficiency in games. In particular, we show that modifying the nonstrategic component of a game properly, it is possible to make all Nash equilibria Pareto optimal. Potential and harmonic games are related to other well-known classes of games, such as the zero-sum games and identical interest games. In Section 2.5.4, we discuss this relation in the context of bimatrix games. As a preview, in Table 2.3, we summarize some of the properties of potential and harmonic games established in the subsequent sections.

2.5.1 Potential Games

Since the seminal paper of Monderer and Shapley (Monderer and Shapley, 1996b), potential games have been an active research topic. The desirable equilibrium properties and structure of these games played a key role in this. In this section we explain the relation of the potential games to the decomposition in Section 2.4 and briefly discuss their properties.

Recall from Definition 2.2.1 that a game is a potential game if and only if there exists some $\phi \in C_0$ such that $Du = \delta_0 \phi$. This condition implies that a game is potential if and only if the associated flow is globally consistent. Thus, it can be seen from the definition of the subspaces and Theorem 2.4.1 that the set of potential games is actually equivalent to $\mathcal{P} \oplus \mathcal{N}$. For future reference, we summarize this result in the following theorem.

Theorem 2.5.1. *The set of potential games is equal to the subspace $\mathcal{P} \oplus \mathcal{N}$.*

Theorem 2.5.1 implies that potential games are games which only have potential and nonstrategic components. Since this set is a subspace, one can consider *projections* onto the set of potential games, i.e., it is possible to find the closest potential game to a given game. We pursue the idea of projection in Section 2.6. Using the previous theorem we next find the dimension of the subspace of potential games (similar results can be found in (Sandholm, 2010b) and (Monderer and Shapley, 1996b)).

Corollary 2.5.1. *The subspace of potential games, $\mathcal{P} \oplus \mathcal{N}$, has dimension $\prod_{m \in \mathcal{M}} h_m + \sum_{m \in \mathcal{M}} \prod_{k \neq m} h_k - 1$.*

Proof. The result immediately follows from Theorem 2.5.1 and Proposition 2.4.1. \square

We next provide a brief discussion of the equilibrium properties of potential games.

Theorem 2.5.2 ((Monderer and Shapley, 1996b)). *Let $\mathcal{G} = \langle \mathcal{M}, \{E^m\}, \{u^m\} \rangle$ be a potential game and ϕ be a corresponding potential function.*

1. *The equilibrium set of \mathcal{G} coincides with the equilibrium set of $\mathcal{G}_\phi \triangleq \langle \mathcal{M}, \{E^m\}, \{\phi\} \rangle$.*
2. *\mathcal{G} has a pure Nash equilibrium.*

The first result follows from the fact that the games \mathcal{G} and \mathcal{G}_ϕ are strategically equivalent. Alternatively, the preferences in \mathcal{G} are aligned with the global objective denoted by the potential function ϕ . The second result is implied by the first one since in finite games the potential function ϕ necessarily has a maximum, and the maximum is a Nash equilibrium of \mathcal{G}_ϕ . These results indicate that potential games can be analyzed by an equivalent game where each player has the same utility function ϕ . The second game is easy to analyze since when agents have the same objective, the game is similar to an optimization problem with objective function ϕ .

2.5.2 Harmonic Games

In this section, we focus on games in which the potential component is zero, hence the strategic interactions are governed only by the harmonic component. We refer to such games as *harmonic games*, i.e., a game \mathcal{G} is a harmonic game if $\mathcal{G} \in \mathcal{H} \oplus \mathcal{N}$. We first provide an alternative definition of harmonic games in terms of the payoff functions of players (cf. equation (2.2)).

Theorem 2.5.3. *The set of games which satisfy*

$$\sum_{m \in \mathcal{M}} \sum_{\mathbf{q}^m \in E^m} (u^m(\mathbf{p}^m, \mathbf{p}^{-m}) - u^m(\mathbf{q}^m, \mathbf{p}^{-m})) = 0,$$

for all strategy profiles \mathbf{p} and strategies $\mathbf{q}^m \in E^m$, is equivalent to the set of harmonic games $\mathcal{H} \oplus \mathcal{N}$.

Proof. The condition in the theorem statement can equivalently be expressed as $\delta_0^* Du = 0$, using the operators introduced in Section 2.4. Since the flow associated with a game is given by Du , a game satisfies this condition if and only if it belongs to $\mathcal{H} \oplus \mathcal{N}$, as can be seen from Definition 2.4.2. \square

This theorem provides a certificate that can be used to check whether a given game is a harmonic game. For instance, using the theorem it immediately follows that the rock-paper-scissors game (for $x = y = z = \frac{1}{3}$) in Example 2.4.1 is a harmonic game.

The rest of this section studies the properties of equilibria of harmonic games. We first characterize the Nash equilibria of such games, and show that generically they do not have a pure Nash equilibrium. We further consider mixed Nash and correlated equilibria, and show how the properties of harmonic games restrict the possible set of equilibria.

Pure Equilibria

In this section, we focus on pure Nash equilibria in harmonic games. Additionally, we characterize the dimension of the space of harmonic games, $\mathcal{H} \oplus \mathcal{N}$.

We first show that at a pure Nash equilibrium of a harmonic game, all players are indifferent between *all* of their strategies.

Lemma 2.5.1. *Let $\mathcal{G} = \langle \mathcal{M}, \{E^m\}, \{u^m\} \rangle$ be a harmonic game and \mathbf{p} be a pure Nash equilibrium. Then,*

$$u^m(\mathbf{p}^m, \mathbf{p}^{-m}) = u^m(\mathbf{q}^m, \mathbf{p}^{-m}) \quad \text{for all } m \in \mathcal{M} \text{ and } \mathbf{q}^m \in E^m. \quad (2.32)$$

Proof. By definition, in harmonic games the utility functions $u = \{u^m\}$ satisfy the condition $\delta_0^* Du = 0$. By (2.13) and (2.21), $\delta_0^* Du$ evaluated at \mathbf{p} can be expressed as,

$$\sum_{m \in \mathcal{M}} \sum_{\mathbf{q} | (\mathbf{p}, \mathbf{q}) \in A^m} (u^m(\mathbf{p}) - u^m(\mathbf{q})) = 0. \quad (2.33)$$

Since \mathbf{p} is a Nash equilibrium it follows that $u^m(\mathbf{p}) - u^m(\mathbf{q}) \geq 0$ for all $(\mathbf{p}, \mathbf{q}) \in A^m$ and $m \in \mathcal{M}$. Combining this with (2.33) it follows that $u^m(\mathbf{p}) - u^m(\mathbf{q}) = 0$ for all $(\mathbf{p}, \mathbf{q}) \in A^m$ and $m \in \mathcal{M}$. Observing that $(\mathbf{p}, \mathbf{q}) \in A^m$ if and only if $\mathbf{q} = (\mathbf{q}^m, \mathbf{p}^{-m})$, the result follows. \square

Using this result we next prove that harmonic games generically do not have pure Nash equilibria. By “generically”, we mean that it is true for almost all harmonic games, except possibly for a set of measure zero (for instance, the trivial game where all utilities are zero is harmonic, and clearly has pure Nash equilibria).

Proposition 2.5.1. *Harmonic games generically do not have pure Nash equilibria.*

Proof. Define $\mathcal{G}_{\mathbf{p}} \subset \mathcal{H} \oplus \mathcal{N}$ as the set of harmonic games for which \mathbf{p} is a pure Nash equilibrium. Observe that $\cup_{\mathbf{p} \in E} \mathcal{G}_{\mathbf{p}}$ is the set of all harmonic games which have a pure Nash equilibria. We show that $\mathcal{G}_{\mathbf{p}}$ is a lower dimensional subspace of the space of harmonic games for each $\mathbf{p} \in E$. Since the set of harmonic games with pure Nash equilibrium is a finite union of lower dimensional subspaces it follows that generically harmonic games do not have pure Nash equilibria.

By Lemma 2.5.1 it follows that

$$\mathcal{G}_{\mathbf{p}} = (\mathcal{H} \oplus \mathcal{N}) \cap \{ \{u^m\}_{m \in \mathcal{M}} \mid u^m(\mathbf{p}) = u^m(\mathbf{q}), \text{ for all } \mathbf{q} \text{ such that } (\mathbf{p}, \mathbf{q}) \in A^m \text{ and } m \in \mathcal{M} \}.$$

Hence $\mathcal{G}_{\mathbf{p}}$ is a subspace contained in $\mathcal{H} \oplus \mathcal{N}$. It immediately follows that $\mathcal{G}_{\mathbf{p}}$ is a lower dimensional subspace if we can show that there exists harmonic games which are not in $\mathcal{G}_{\mathbf{p}}$, i.e., in which \mathbf{p} is not a pure Nash equilibrium.

Assume that \mathbf{p} is a pure Nash equilibrium in all harmonic games. Since \mathbf{p} is arbitrary this holds only if all strategy profiles are pure Nash equilibria in harmonic games. If all strategy profiles are Nash equilibria, by Lemma 2.5.1 it follows that the pairwise ranking function is equal to zero in harmonic games, hence $\mathcal{H} \oplus \mathcal{N} \subset \mathcal{N}$. We reach a contradiction since dimension of \mathcal{H} is larger than zero.

Therefore, $\mathcal{G}_{\mathbf{p}}$ is a strict subspace of the space of harmonic games, and thus harmonic games generically do not have pure Nash equilibria. \square

We conclude this section by a dimension result that is analogous to the result obtained for potential games.

Theorem 2.5.4. *The set of harmonic games, $\mathcal{H} \oplus \mathcal{N}$, has dimension $(M-1) \prod_{m \in \mathcal{M}} h_m + 1$.*

Proof. The result immediately follows from Theorem 2.4.1 and Proposition 2.4.1. \square

Mixed Nash and Correlated Equilibria in Harmonic Games

In the previous section we showed that harmonic games generically do not have pure Nash equilibria. In this section, we study their mixed Nash and correlated equilibria. In particular, we show that in harmonic games, the mixed strategy profile, in which players uniformly randomize over their strategies is always a mixed Nash equilibrium. Additionally, in the case of two-player harmonic games mixed Nash and correlated equilibria coincide, and if players have equal number of strategies the uniformly mixed strategy profile is the unique correlated equilibrium of the game. Before we discuss the details of these results, we next provide some preliminaries and notation.

We denote the set of probability distributions on E by ΔE . Given $x \in \Delta E$, $x(\mathbf{p})$ denotes the probability assigned to $\mathbf{p} \in E$. Observe that for all $x \in \Delta E$, $\sum_{\mathbf{p} \in E} x(\mathbf{p}) = 1$, and $x(\mathbf{p}) \geq 0$. Similarly for each player $m \in \mathcal{M}$, ΔE^m denotes the set of probability distributions on E^m and for $x^m \in \Delta E^m$, $x^m(\mathbf{p}^m)$ is the probability assigned to strategy $\mathbf{p}^m \in E^m$. As before all $x^m \in \Delta E^m$ satisfies $\sum_{\mathbf{p}^m \in E^m} x^m(\mathbf{p}^m) = 1$ and $x^m(\mathbf{p}^m) \geq 0$. We refer to the distribution $x^m \in \Delta E^m$ as a mixed strategy of player $m \in \mathcal{M}$ and the collection $x = \{x^m\}_m$ as a mixed strategy profile. Note that $\{x^m\}_m \in \prod_{m \in \mathcal{M}} \Delta E^m \subset \Delta E$. Mixed strategies of all players but the m th one is denoted by x^{-m} .

With some abuse of the notation, we define the mixed extensions of the utility functions

$u^m : \prod_{m \in \mathcal{M}} \Delta E^m \rightarrow \mathbb{R}$ such that for any $x \in \prod_{m \in \mathcal{M}} \Delta E^m$,

$$u^m(x) = \sum_{\mathbf{p} \in E} u^m(\mathbf{p}) \prod_{k \in \mathcal{M}} x^k(\mathbf{p}^k). \quad (2.34)$$

Similarly, if player m uses pure strategy \mathbf{q}^m and the other players use the mixed strategies x^{-m} we denote the payoff of player m by,

$$u^m(\mathbf{q}^m, x^{-m}) = \sum_{\mathbf{p}^{-m} \in E^{-m}} u^m(\mathbf{q}^m, \mathbf{p}^{-m}) \prod_{k \in \mathcal{M}, k \neq m} x^k(\mathbf{p}^k). \quad (2.35)$$

Using this notation we can define the solution concepts.

Definition 2.5.1 (Mixed Nash / Correlated Equilibrium). *Consider the game $\langle \mathcal{M}, \{E^m\}, \{u^m\} \rangle$.*

1. *A mixed strategy profile $x = \{x^m\}_m \in \prod_{m \in \mathcal{M}} \Delta E^m$ is a mixed Nash equilibrium if for all $m \in \mathcal{M}$ and $\mathbf{p}^m \in E^m$, $u^m(x^m, x^{-m}) \geq u^m(\mathbf{p}^m, x^{-m})$.*
2. *A probability distribution $x \in \Delta E$ is a correlated equilibrium if for all $m \in \mathcal{M}$ and $\mathbf{p}^m, \mathbf{q}^m \in E^m$, $\sum_{\mathbf{p}^{-m} \in E^{-m}} (u^m(\mathbf{p}^m, \mathbf{p}^{-m}) - u^m(\mathbf{q}^m, \mathbf{p}^{-m})) x(\mathbf{p}^m, \mathbf{p}^{-m}) \geq 0$.*

From these definitions it can be seen that every mixed Nash equilibrium is a correlated equilibrium where the corresponding distribution $x \in \prod_{m \in \mathcal{M}} \Delta E^m \subset \Delta E$ is a product distribution, i.e., it satisfies $x(\mathbf{p}) = \prod_{m \in \mathcal{M}} x^m(\mathbf{p}^m)$

These definitions also imply that similar to Nash equilibrium, the conditions for mixed Nash and correlated equilibria can be expressed only in terms of pairwise comparisons. Therefore, these equilibrium sets are independent of the nonstrategic components of games.

We next obtain an alternative characterization of correlated equilibria in normalized harmonic games. This characterization will be more convenient when studying the equilibrium properties of harmonic games, as it is expressed in terms of equalities, instead of inequalities. See Section 2.8 for a proof.

Proposition 2.5.2. *Consider a normalized harmonic game, $\mathcal{G} = \langle \mathcal{M}, \{u^m\}, E^m \rangle$ and a probability distribution $x \in \Delta E$. The following are equivalent:*

1. *x is a correlated equilibrium.*
2. *For all $\mathbf{p}^m, \mathbf{q}^m$ and $m \in \mathcal{M}$,*

$$\sum_{\mathbf{p}^{-m} \in E^{-m}} (u^m(\mathbf{p}^m, \mathbf{p}^{-m}) - u^m(\mathbf{q}^m, \mathbf{p}^{-m})) x(\mathbf{p}^m, \mathbf{p}^{-m}) = 0. \quad (2.36)$$

3. *For all $\mathbf{p}^m, \mathbf{q}^m$ and $m \in \mathcal{M}$,*

$$\sum_{\mathbf{p}^{-m} \in E^{-m}} u^m(\mathbf{q}^m, \mathbf{p}^{-m}) x(\mathbf{p}^m, \mathbf{p}^{-m}) = 0. \quad (2.37)$$

The above proposition implies that the correlated equilibria of harmonic games correspond to the intersection of the probability simplex with a subspace defined by the utilities in the game. Using this result, we obtain the following characterization of mixed Nash equilibria of harmonic games.

Corollary 2.5.2. *Let $\mathcal{G} = \langle \mathcal{M}, \{u^m\}, \{E^m\} \rangle$ be a harmonic game. The mixed strategy profile $x \in \prod_{m \in \mathcal{M}} \Delta E^m$ is a mixed Nash equilibrium if and only if,*

$$u^m(x^m, x^{-m}) = u^m(\mathbf{p}^m, x^{-m}) \quad \text{for all } \mathbf{p}^m \in E^m \text{ and } m \in \mathcal{M}. \quad (2.38)$$

Proof. Assume that (2.38) holds, then clearly all players are indifferent between all their mixed strategies, hence it follows that x is a mixed Nash equilibrium of the game.

Let x be a mixed Nash equilibrium. Since each mixed Nash equilibrium is also a correlated equilibrium, from equivalence of (i) and (ii) of Proposition 2.5.2 for all harmonic games, it follows that for all $\mathbf{p}^m, \mathbf{q}^m$ and $m \in \mathcal{M}$,

$$\begin{aligned} 0 &= \sum_{\mathbf{p}^{-m} \in E^{-m}} (u^m(\mathbf{p}^m, \mathbf{p}^{-m}) - u^m(\mathbf{q}^m, \mathbf{p}^{-m})) x(\mathbf{p}^m, \mathbf{p}^{-m}) \\ &= x^m(\mathbf{p}^m) \sum_{\mathbf{p}^{-m} \in E^{-m}} (u^m(\mathbf{p}^m, \mathbf{p}^{-m}) - u^m(\mathbf{q}^m, \mathbf{p}^{-m})) \prod_{k \neq m} x^k(\mathbf{p}^k) \\ &= x^m(\mathbf{p}^m) (u^m(\mathbf{p}^m, x^{-m}) - u^m(\mathbf{q}^m, x^{-m})). \end{aligned} \quad (2.39)$$

Since by definition of probability distributions, there exists \mathbf{p}^m such that $x^m(\mathbf{p}^m) > 0$ it follows that $u^m(\mathbf{p}^m, x^{-m}) = u^m(\mathbf{q}^m, x^{-m})$ for all $\mathbf{q}^m \in E^m$. Thus, $u^m(x^m, x^{-m}) = u^m(\mathbf{q}^m, x^{-m})$ for all $\mathbf{q}^m \in E^m$. Since m is arbitrary, the claim follows. \square

It is well-known that in mixed Nash equilibria of games, players are indifferent between all the pure strategies in the support of their mixed strategy (see (Fudenberg and Tirole, 1991)), i.e., if $x \in \prod_{m \in \mathcal{M}} \Delta E^m$ is a mixed Nash equilibrium then

$$u^m(x^m, x^{-m}) \begin{cases} = u^m(\mathbf{p}^m, x^{-m}) & \text{for all } \mathbf{p}^m \text{ such that } x^m(\mathbf{p}^m) \geq 0 \\ \geq u^m(\mathbf{p}^m, x^{-m}) & \text{for all } \mathbf{p}^m \text{ such that } x^m(\mathbf{p}^m) = 0. \end{cases} \quad (2.40)$$

The above corollary implies that at a mixed equilibrium of a harmonic game, each player is indifferent between all its pure strategies, including those which are not in the support of its mixed strategy.

We next define a particular mixed strategy profile, and show that it is an equilibrium in all harmonic games.

Definition 2.5.2 (Uniformly Mixed Strategy Profile). *The uniformly mixed strategy of player m is a mixed strategy where player m uses $x_{\mathbf{q}^m} = \frac{1}{h_m}$ for all $\mathbf{q}^m \in E^m$. Respectively, we define the uniformly mixed strategy profile as the one in which all players use uniformly mixed strategies.*

Recall that rock-paper-scissors and matching pennies are examples of harmonic games, in which the uniformly mixed strategy profile is a mixed Nash equilibrium. The next theorem shows that this is a general property of harmonic games and the uniformly mixed strategy profile is always a Nash equilibrium.

Theorem 2.5.5. *In harmonic games, the uniformly mixed strategy profile is always a Nash equilibrium.*

Proof. Let $\mathcal{G} = \langle \mathcal{M}, \{u^m\}, \{E^m\} \rangle$ be a harmonic game, and x be the uniformly mixed strategy profile. In order to prove the claim we first state the following useful identity (see Section 2.8 for a proof), on the utility functions of harmonic games.

Lemma 2.5.2. *Let $\mathcal{G} = \langle \mathcal{M}, \{u^m\}, \{E^m\} \rangle$ be a harmonic game. Then for all $\mathbf{q}^m, \mathbf{r}^m \in E^m$, $m \in \mathcal{M}$, $\sum_{\mathbf{p}^{-m} \in E^{-m}} u^m(\mathbf{r}^m, \mathbf{p}^{-m}) - u^m(\mathbf{q}^m, \mathbf{p}^{-m}) = 0$.*

Using this lemma, and the definition of the uniformly mixed strategy, it follows that

$$\begin{aligned} u^m(\mathbf{q}^m, x^{-m}) - u^m(\mathbf{p}^m, x^{-m}) &= \sum_{\mathbf{p}^{-m} \in E^{-m}} c_m (u^m(\mathbf{q}^m, \mathbf{p}^{-m}) - u^m(\mathbf{p}^m, \mathbf{p}^{-m})) \\ &= c_m \sum_{\mathbf{p}^{-m} \in E^{-m}} (u^m(\mathbf{q}^m, \mathbf{p}^{-m}) - u^m(\mathbf{p}^m, \mathbf{p}^{-m})) \quad (2.41) \\ &= 0, \end{aligned}$$

where $c_m = \prod_{k \neq m} x^k(\mathbf{p}^k) = \prod_{k \neq m} \frac{1}{h_k}$. Since \mathbf{p}^m and \mathbf{q}^m are arbitrary, (2.41) implies that

$$u^m(x^m, x^{-m}) = u^m(\mathbf{p}^m, x^{-m}) \quad (2.42)$$

for all $\mathbf{p}^m \in E^m$, and by Corollary 2.5.2, x is a mixed strategy Nash equilibrium. \square

We conclude this section, by providing an additional characterization of normalized harmonic games that is useful for the study of the equilibrium properties of these games.

Theorem 2.5.6. *The game \mathcal{G} with utilities $u = \{u^m\}_{m \in \mathcal{M}}$ is a normalized harmonic game, i.e., it belongs to \mathcal{H} if and only if $\sum_{m \in \mathcal{M}} h_m u^m = 0$ and $\Pi_m u^m = u^m$ for all $m \in \mathcal{M}$, where $h_m = |E^m|$.*

Proof. By Definition 2.4.2, $\mathcal{G} \in \mathcal{H}$ if and only if $\Pi u = u$ and $\delta_0^* D u = 0$. Using the definitions of the operators, these conditions can alternatively be expressed as $\Pi_m u^m = u^m$ and $\delta_0^* \sum_{m \in \mathcal{M}} D_m u^m = 0$. By (2.24) and the orthogonality of image spaces of operators D_m , the latter equality implies that $\sum_{m \in \mathcal{M}} D_m^* D_m u^m = \sum_{m \in \mathcal{M}} \Delta_{0,m} u^m = 0$. Using Lemma 2.4.1, $\Delta_{0,m} = h_m \Pi_m$, and hence it follows that $\mathcal{G} \in \mathcal{H}$, if and only if

$$\sum_{m \in \mathcal{M}} h_m \Pi_m u^m = 0 \quad \text{and,} \quad \Pi_m u^m = u^m \text{ for all } m. \quad (2.43)$$

The claim follows by replacing $\Pi_m u^m$ in the summation with u^m . \square

The above theorem implies that normalized harmonic games, where players have an equal number of strategies, are *zero-sum games*, i.e., in such games the payoffs of players add up to zero at all strategy profiles. This suggests that equilibrium properties of zero-sum games and harmonic games should be related. We study the relation between zero-sum games and harmonic games in Section 2.5.4.

For two player harmonic games, using the above zero-sum characterization, additional properties of equilibrium can be established. For instance, in these games sets of correlated equilibria and mixed equilibria are identical. These properties are explored in Section 2.9.

2.5.3 Nonstrategic Component and Efficiency in Games

As discussed in Section 2.4.2, the pairwise comparisons in a game are functions of only the potential and harmonic components of the game. Thus, the nonstrategic component has no effect on the equilibrium properties of games. However, the nonstrategic component is of interest mainly through its effect on the efficiency properties of games, as discussed in the rest of this section. The efficiency measure we focus on is Pareto optimality.

Definition 2.5.3 (Pareto Optimality). *A strategy profile \mathbf{p} is Pareto optimal if and only if there does not exist another strategy profile \mathbf{q} such that all players weakly increase their payoffs and one player strictly increases its payoff, i.e.,*

$$\begin{aligned} u^m(\mathbf{q}) &\geq u^m(\mathbf{p}), & \text{for all } m \in \mathcal{M} \\ u^k(\mathbf{q}) &> u^k(\mathbf{p}), & \text{for some } k \in \mathcal{M}. \end{aligned} \tag{2.44}$$

We first state a preliminary lemma, which will be useful in the subsequent analysis.

Lemma 2.5.3. *Let \mathcal{G} be a game with utilities $\{u^m\}$. There exists a game $\hat{\mathcal{G}}$ with utilities $\{\hat{u}^m\}$ such that (i) the potential and harmonic components of $\hat{\mathcal{G}}$ are identical to those of \mathcal{G} and (ii) in $\hat{\mathcal{G}}$ all players get zero payoff at all strategy profiles that are pure Nash equilibria of \mathcal{G} .*

Proof. Let $N_{\mathcal{G}}$ be the set of pure Nash equilibria of \mathcal{G} . First observe that if there are m -comparable equilibria in \mathcal{G} player m receives the same payoff in these equilibria, i.e., if $\mathbf{p}, \mathbf{q} \in N_{\mathcal{G}}$ and $\mathbf{p} = (\mathbf{p}^m, \mathbf{p}^{-m})$, $\mathbf{q} = (\mathbf{q}^m, \mathbf{p}^{-m})$ for some m , then $u^m(\mathbf{p}^m, \mathbf{p}^{-m}) = u^m(\mathbf{q}^m, \mathbf{p}^{-m})$. This equality holds since otherwise, player m would have incentive to improve its payoff at \mathbf{p} or \mathbf{q} by switching to a strategy profile with better payoff, and this contradicts with \mathbf{p} and \mathbf{q} being Nash equilibria of \mathcal{G} .

Define the game $\hat{\mathcal{G}}$ with utilities $\{\hat{u}^m\}_{m \in \mathcal{M}}$ such that

$$\hat{u}^m(\mathbf{p}) = \begin{cases} 0 & \mathbf{p} \in N_{\mathcal{G}} \\ u^m(\mathbf{p}) - u^m(\mathbf{q}) & \text{if there exists a } \mathbf{q} \in N_{\mathcal{G}} \text{ which is } m\text{-comparable with } \mathbf{p} \\ u^m(\mathbf{p}) & \text{otherwise.} \end{cases} \tag{2.45}$$

for all $m \in \mathcal{M}$, $\mathbf{p} \in E$. Note that \hat{u}^m is well defined since in \mathcal{G} player m gets the same payoff in all $\mathbf{p} \in N_{\mathcal{G}}$ that are m -comparable. Note that in $\hat{\mathcal{G}}$ all players receive zero payoff at all strategy profiles $\mathbf{p} \in N_{\mathcal{G}}$. To prove the claim, it suffices to show that \mathcal{G} and $\hat{\mathcal{G}}$ have the same potential and harmonic components, or equivalently the game with utilities $\{u^m - \hat{u}^m\}_{m \in \mathcal{M}}$ is nonstrategic, i.e., belongs to \mathcal{N} .

In order to prove that the difference is nonstrategic, we first show that the pairwise comparisons of games with utilities $\{u^m\}_{m \in \mathcal{M}}$ and $\{\hat{u}^m\}_{m \in \mathcal{M}}$ are the same. Note that by (2.45) given m -comparable \mathbf{p} and \mathbf{q} , $u^m(\mathbf{p}) - u^m(\mathbf{q}) = \hat{u}^m(\mathbf{p}) - \hat{u}^m(\mathbf{q})$, if there is no $\mathbf{r} \in N_{\mathcal{G}}$ that is m -comparable with \mathbf{p} or \mathbf{q} . If there exists $\mathbf{r} \in N_{\mathcal{G}}$ that is m -comparable with \mathbf{p} , then it is also m -comparable with \mathbf{q} , hence it follows by (2.45) that $\hat{u}^m(\mathbf{p}) - \hat{u}^m(\mathbf{q}) = u^m(\mathbf{p}) - u^m(\mathbf{r}) - u^m(\mathbf{q}) + u^m(\mathbf{r}) = u^m(\mathbf{p}) - u^m(\mathbf{q})$. Note that these equalities hold even if \mathbf{p} or \mathbf{q} is in $N_{\mathcal{G}}$.

Thus, for any m -comparable \mathbf{p} and \mathbf{q} it follows that

$$(u^m(\mathbf{p}) - \hat{u}^m(\mathbf{p})) - (u^m(\mathbf{q}) - \hat{u}^m(\mathbf{q})) = 0,$$

hence the game with utilities $\{u^m - \hat{u}^m\}_{m \in \mathcal{M}}$ is nonstrategic and the claim follows. \square

Note that if two games differ only in their nonstrategic components, the pairwise comparisons, and hence the equilibria of these games are identical. Therefore, an immediate implication of the above lemma is that for a given game there exists another game with same equilibrium set such that the payoffs at all Nash equilibria are equal to zero. We use this to prove the following Pareto optimality result.

Theorem 2.5.7. *Let \mathcal{G} be a game with utilities $\{u^m\}$. There exists a game $\bar{\mathcal{G}}$ with utilities $\{\bar{u}^m\}$ such that (i) the potential and harmonic components of $\bar{\mathcal{G}}$ are identical to these of \mathcal{G} and (ii) in $\bar{\mathcal{G}}$ the set of pure NE coincides with the set of Pareto optimal strategy profiles.*

Proof. Games that differ only in nonstrategic components have identical pairwise comparisons, hence the set of Nash equilibria (NE) is the same for such games. Let $N_{\mathcal{G}}$ denote the set of pure NE of \mathcal{G} , or equivalently the set of pure NE of a game which differs from \mathcal{G} only by its nonstrategic component.

By Lemma 2.5.3, it follows that for any game \mathcal{G} there exists a game such that the two games differ only in their nonstrategic components and all players receive zero payoffs at all pure NE (strategy profiles in $N_{\mathcal{G}}$). Therefore, without loss of generality, we let \mathcal{G} be a game in which all players receive zero payoffs at all NE. Given such a game, let $\alpha = 1 + \max_{m, \mathbf{p}} u^m(\mathbf{p})$. Consider the game $\bar{\mathcal{G}}$ with utilities $\{\bar{u}^m\}_{m \in \mathcal{M}}$ such that

$$\bar{u}^m(\mathbf{p}) = \begin{cases} u^m(\mathbf{p}) & \text{if } \mathbf{p} \in N_{\mathcal{G}} \text{ or if there exists a } \mathbf{q} \in N_{\mathcal{G}} \text{ which is } m\text{-comparable with } \mathbf{p} \\ u^m(\mathbf{p}) - \alpha & \text{otherwise.} \end{cases}$$

for all $m \in \mathcal{M}$, $\mathbf{p} \in E$.

Consider m -comparable strategy profiles \mathbf{p} and \mathbf{q} . Observe that if there exists a strategy profile \mathbf{r} that is m -comparable with \mathbf{p} , it is also m -comparable with \mathbf{q} since by definition of m -comparable strategy profiles $\mathbf{p}^{-m} = \mathbf{r}^{-m} = \mathbf{q}^{-m}$.

Assume that there is a NE that is m -comparable with \mathbf{p} or \mathbf{q} , then by definition of \bar{u}^m it follows that $u^m(\mathbf{p}) - u^m(\mathbf{q}) = \bar{u}^m(\mathbf{p}) - \bar{u}^m(\mathbf{q})$. On the contrary if there is no NE that is m -comparable with \mathbf{p} or \mathbf{q} then $\bar{u}^m(\mathbf{p}) - \bar{u}^m(\mathbf{q}) = u^m(\mathbf{p}) - \alpha - u^m(\mathbf{q}) + \alpha = u^m(\mathbf{p}) - u^m(\mathbf{q})$. Hence \mathcal{G} and $\bar{\mathcal{G}}$ have identical pairwise comparisons, and thus the game with utilities $\{u^m - \bar{u}^m\}_{m \in \mathcal{M}}$ is nonstrategic.

We prove the claim, by showing that at all strategy profiles that are not an equilibrium in $\bar{\mathcal{G}}$ (equivalently in \mathcal{G}), the players receive nonpositive payoffs and at least one player receives negative payoff and at all NE all players receive zero payoff. This immediately implies that strategy profiles, that are not NE cannot be Pareto optimal, as deviation to a NE increases the payoff of at least one player and the payoff of other players do not decrease by such a deviation. Additionally, it implies that all NE are Pareto optimal, since at all NE all players receive the same payoff, and deviation to a strategy profile that is not a NE strictly decreases the payoff of at least a single player.

By construction it follows that at all NE all players receive zero payoff. Let \mathbf{p} be a strategy profile that is not a NE. If there is some m for which \mathbf{p} is not m -comparable to a NE, then it follows that $\bar{u}^m(\mathbf{p}) = u^m(\mathbf{p}) - \alpha \leq -1$. If on the other hand, \mathbf{p} is m -comparable to a NE, then $\bar{u}^m(\mathbf{p}) \leq 0$, since payoffs are equal to zero at NE. Thus, at any strategy profile, \mathbf{p} , that is not a NE players receive nonpositive payoffs, and additionally if for some player m , \mathbf{p} is not m -comparable to a NE, player m receives strictly negative payoff.

To finish the proof we need to show that if \mathbf{p} is m -comparable to a NE for all $m \in \mathcal{M}$, then it still follows that $\bar{u}^m(\mathbf{p}) < 0$ for some $m \in \mathcal{M}$. Assume that this is not true and $\bar{u}^m(\mathbf{p}) = 0$ for all $m \in \mathcal{M}$. Since \mathbf{p} is not a NE, there is at least one player, say m , who can get strictly positive payoff by deviating to a different strategy profile. Therefore this player has strictly positive payoff after its deviation. However, as argued earlier payoffs are nonpositive at strategy profiles that are not NE, and zero at NE. Thus. we reach a contradiction and $\bar{u}^m(\mathbf{p}) < 0$ for some $m \in \mathcal{M}$.

Therefore, it follows that all players have zero payoffs at all NE, and at any other strategy profile all players have nonpositive payoffs and at least one player has strictly negative payoff. \square

Note that it is possible to obtain similar results for other efficiency measures using similar arguments to those given in this section. This direction will not be pursued in this chapter. The above theorem suggests that the difference in the nonstrategic component of games that are otherwise identical may cause the efficiency properties of these game to be very different. In particular, in one of the games all equilibria may be Pareto optimal when this is not the case for the other game. Therefore, although the nonstrategic component does not change the pairwise comparisons and equilibrium properties in a game it plays a

key role in Pareto optimality of equilibria.

2.5.4 Zero-Sum Games and Identical Interest Games

In this section we present a different decomposition of the space of games, and discuss its relation to our decomposition. To simplify the presentation, we focus on bimatrix games, where each player has h strategies. Before introducing the decomposition, we define zero-sum games and identical interest games.

Definition 2.5.4 (Zero-sum and Identical Interest Games). *Let \mathcal{G} denote the bimatrix game with payoff matrices (A, B) . \mathcal{G} is a zero-sum game, if $A + B = 0$, and \mathcal{G} is an identical interest game, if $A = B$.*

We denote the set of zero-sum games by \mathcal{Z} , and the set of identical interest games by \mathcal{I} . Since these sets are defined by equality constraints on the payoff matrices, it follows that they are subspaces.

The idea of decomposing a game to an identical interest game and a zero-sum game was previously mentioned in the literature for two-player games, (Başar and Ho, 1974). The following lemma implies that \mathcal{Z} and \mathcal{I} decomposition of the set of games, has the direct sum property.

Lemma 2.5.4. *The space of two-player games $\mathcal{G}_{\mathcal{M},E}$ is a direct sum of subspaces of zero-sum and identical interest games, i.e., $\mathcal{G}_{\mathcal{M},E} = \mathcal{Z} \oplus \mathcal{I}$.*

Proof. Consider a bimatrix game with utilities (u^1, u^2) . Observe that this game can be decomposed to the games with payoff functions $(\frac{u^1 - u^2}{2}, \frac{u^2 - u^1}{2})$ and $(\frac{u^1 + u^2}{2}, \frac{u^1 + u^2}{2})$. Clearly the former game is a zero-sum game, where the latter is an identical interest game. Since the initial game was arbitrary, it follows that any game can be decomposed to a zero-sum game and an identical interest game. The direct sum property follows, since for two-player zero-sum and identical interest games, with utility functions $(u, -u)$ and (v, v) respectively, if $(u + v, u - v) = (0, 0)$, then $u = v = 0$. \square

Note that Theorem 2.5.6 suggests that two-player normalized harmonic games, where players have equal number of strategies are zero-sum.⁷ Also, it immediately follows by checking the definitions that identical interest games are potential games. This intuitively suggests that the zero-sum and identical interest game decomposition closely relates to our decomposition. In the following theorem, we establish this relation by characterizing the dimensions of the intersections of the subspaces \mathcal{Z} and \mathcal{I} , with the sets of potential and harmonic games. We provide a proof in Section 2.8.

⁷In addition, if the definition of zero-sum is generalized to include multiplayer games where payoffs of all players add up to zero, then it can be seen that normalized harmonic games where players have equal number of strategies are still zero-sum games.

Theorem 2.5.8. Consider two-player games, in which each player has h strategies. The dimensions of intersections of the subspaces of zero-sum and identical interest games (\mathcal{Z} and \mathcal{I}) with the subspaces of potential and harmonic games ($\mathcal{P} \oplus \mathcal{N}$ and $\mathcal{H} \oplus \mathcal{N}$) are as in the following table.

	\mathcal{Z}	\mathcal{I}	$\mathcal{Z} \oplus \mathcal{I}$
$\mathcal{P} \oplus \mathcal{N}$	$2h - 1$	h^2	$h^2 + 2h - 1$
$\mathcal{H} \oplus \mathcal{N}$	$h^2 - 2h + 2$	1	$h^2 + 1$
$\mathcal{P} \oplus \mathcal{H} \oplus \mathcal{N}$	h^2	h^2	$2h^2$

Table 2.4: Dimensions of subspaces of games and their intersections

The above theorem suggests that the dimensions of harmonic games and zero-sum games (and similarly identical interest games and potential games) are close to the dimension of their intersections. Thus, zero-sum games are in general closely related to harmonic games, and identical interest games are related to potential games. On the other hand, it is possible to find instances of zero-sum games that are potential games, and not harmonic games (see Table 2.5).

	a	b
x	0, 0	1, -1
y	-1, 1	0, 0

(a) Payoffs

	a	b
x	2	1
y	1	0

(b) Potential function

Table 2.5: A zero-sum potential game

In general, the identical interest component is a potential game, and it can be used to approximate a given game with a potential game. However, as illustrated in Table 2.6, this approximation need not yield the closest potential game to a given game. In this example, despite the fact that the original game is a potential game, the zero-sum and identical interest game decomposition may lead to a potential game which is much farther than the closest potential game

We believe that the decomposition presented in Section 2.4 is more natural than the zero-sum identical interest game decomposition, as it clearly separates the strategic ($\mathcal{P} \oplus \mathcal{H}$) and nonstrategic (\mathcal{N}) components of games and further identifies components, such as potential and harmonic components, with distinct strategic properties. In addition, it is invariant under trivial manipulations that do not change the strategic interactions, i.e., changes in the nonstrategic component.

	a	b
x	1, 1	1, -1
y	-1, 1	-1, -1

(a) Payoffs in \mathcal{G}

	a	b
x	4	2
y	2	0

(b) Potential function of \mathcal{G}

	a	b
x	0, 0	1, -1
y	-1, 1	0, 0

(c) Payoffs in \mathcal{G}_Z

	a	b
x	1, 1	0, 0
y	0, 0	-1, -1

(d) Payoffs in \mathcal{G}_I

Table 2.6: A potential game \mathcal{G} and its zero-sum (\mathcal{G}_Z) and identical interest components (\mathcal{G}_I).

2.6 Projections onto Potential and Harmonic Games

In this section, we discuss projections of games onto the subspaces of potential and harmonic games. In Section 2.4.2, we defined the subspaces $\mathcal{P}, \mathcal{H}, \mathcal{N}$ of potential, harmonic, and nonstrategic components, respectively. We also proved that they provide a direct sum decomposition of the space of all games. In this section, we show that under an appropriately defined inner product in $\mathcal{G}_{\mathcal{M}, E}$, the harmonic, potential and nonstrategic subspaces become *orthogonal*. We use our decomposition result together with this inner product to obtain projections of games to these subspaces, i.e., for an arbitrary game, we present closed-form expressions for the “closest” potential and harmonic games with respect to this inner product.

Let $\mathcal{G}, \hat{\mathcal{G}}$ be two games in $\mathcal{G}_{\mathcal{M}, E}$. We define the inner product on $\mathcal{G}_{\mathcal{M}, E}$ as

$$\langle \mathcal{G}, \hat{\mathcal{G}} \rangle_{\mathcal{M}, E} \triangleq \sum_{m \in \mathcal{M}} h_m \langle u^m, \hat{u}^m \rangle, \quad (2.46)$$

where the inner product in the right hand side is the inner product of C_0 as defined in (2.8), i.e., it is the inner product of the space of functions defined on E . Note that it can be easily checked that (2.46) is an inner product, by observing that it is a weighted version of the standard inner product in C_0^M . The given inner product also induces a norm which will help us quantify the distance between games. We define the norm on $\mathcal{G}_{\mathcal{M}, E}$ as follows:

$$\|\mathcal{G}\|_{\mathcal{M}, E}^2 = \langle \mathcal{G}, \mathcal{G} \rangle_{\mathcal{M}, E}. \quad (2.47)$$

Note that this norm also corresponds to a weighted l_2 norm defined on the space C_0^M .

Next we prove that the potential, harmonic and nonstrategic subspaces are orthogonal under this inner product.

Theorem 2.6.1. *Under the inner product introduced in (2.46), we have $\mathcal{P} \perp \mathcal{H} \perp \mathcal{N}$, i.e., the potential, harmonic and nonstrategic subspaces are orthogonal.*

Proof. Let $\{u_P^m\}_{m \in \mathcal{M}} = \mathcal{G}_P \in \mathcal{P}$, $\{u_H^m\}_{m \in \mathcal{M}} = \mathcal{G}_H \in \mathcal{H}$ and $\{u_N^m\}_{m \in \mathcal{M}} = \mathcal{G}_N \in \mathcal{N}$ be arbitrary games in \mathcal{P} , \mathcal{H} and \mathcal{N} respectively. In order to prove the claim we will first prove $\mathcal{G}_N \perp \mathcal{G}_H$ and $\mathcal{G}_N \perp \mathcal{G}_P$. Secondly we prove $\mathcal{G}_P \perp \mathcal{G}_H$. Since the games are arbitrary the first part will imply that $\mathcal{N} \perp \mathcal{H}$ and $\mathcal{N} \perp \mathcal{P}$ and the second part will imply that $\mathcal{P} \perp \mathcal{H}$ proving the claim.

Note that by definition $u_N^m \in \ker(D_m)$ for all $m \in \mathcal{M}$ and u_P^m, u_H^m are in the orthogonal complement of $\ker(D_m)$ since $\Pi_m u_P^m = u_P^m$, $\Pi_m u_H^m = u_H^m$ and Π_m is the projection operator to the orthogonal complement of $\ker(D_m)$. This implies that $\langle u_P^m, u_N^m \rangle = \langle u_H^m, u_N^m \rangle = 0$ for all $m \in \mathcal{M}$ and hence using the inner product introduced in (2.46) it follows that $\mathcal{G}_N \perp \mathcal{G}_H$ and $\mathcal{G}_N \perp \mathcal{G}_P$.

Next observe that for all $m \in \mathcal{M}$,

$$\langle u_P^m, u_H^m \rangle = \langle D_m^\dagger D_m u_P^m, u_H^m \rangle = \frac{1}{h_m} \langle D_m^* D_m \phi, u_H^m \rangle = \frac{1}{h_m} \langle \phi, D_m^* D_m u_H^m \rangle,$$

where the first equality follows from $\Pi_m u_P^m = u_P^m$, and the second equality follows from Lemma 2.4.1 and the fact that $D_m u_P^m = D_m \phi$. The third equality uses the properties of the operators D_m and D_m^* . Therefore,

$$\langle \mathcal{G}_P, \mathcal{G}_H \rangle_{\mathcal{M}, E} = \sum_{m \in \mathcal{M}} \langle \phi, D_m^* D_m u_H^m \rangle = \langle \phi, \sum_{m \in \mathcal{M}} D_m^* D_m u_H^m \rangle = \langle \phi, \delta_0^* \sum_{m \in \mathcal{M}} D_m u_H^m \rangle = 0.$$

Since $\delta_0^* \sum_{m \in \mathcal{M}} D_m u_H^m = 0$ by the definition of \mathcal{H} . Here the last equality follows using $\delta_0^* = \sum_m D_m^*$ and orthogonality of the image spaces of D_m for $m \in \mathcal{M}$. Therefore, $\mathcal{G}_H \perp \mathcal{G}_P$ as claimed and the result follows. \square

The next theorem provides closed form expressions for the closest potential and harmonic games with respect to the norm in (2.47).

Theorem 2.6.2. *Let $\mathcal{G} \in \mathcal{G}_{\mathcal{M}, E}$ be a game with utilities $\{u^m\}_{m \in \mathcal{M}}$, and let $\phi = \delta_0^\dagger Du$. With respect to the norm in (2.47),*

1. *The closest potential game to \mathcal{G} has utilities $\Pi_m \phi + (I - \Pi_m)u^m$ for all $m \in \mathcal{M}$,*
2. *The closest harmonic game to \mathcal{G} has utilities $u^m - \Pi_m \phi$ for all $m \in \mathcal{M}$.*

Proof. By Theorem 2.6.1, the harmonic component of \mathcal{G} is orthogonal to the space of potential games $\mathcal{P} \oplus \mathcal{N}$. Thus, the closest potential game to \mathcal{G} has utilities $u^m - u_H^m$, where $\{u_H^m\}_{m \in \mathcal{M}}$ is the harmonic component of \mathcal{G} . Similarly, the potential component of \mathcal{G} is orthogonal to the space of harmonic games $\mathcal{H} \oplus \mathcal{N}$ and thus the closest harmonic game to \mathcal{G} has utilities $u^m - u_P^m$, where $\{u_P^m\}_{m \in \mathcal{M}}$ is the potential component of \mathcal{G} . Using the closed form expressions for u_P^m and u_H^m from Theorem 2.4.1, the claim follows. \square

Note that the utilities in the closest potential game consist of two parts: the term $\Pi_m \phi$ expresses the preferences that are captured by the potential function ϕ , and $(I - \Pi_m)u^m$

corresponds to the nonstrategic component of the original game. Similarly, the closest harmonic game differs from the original game by its potential component, and hence has the same nonstrategic and harmonic components with the original game. This implies that the projection decomposes the flows generated by a game to its consistent and inconsistent components and is closely related to the decomposition of flows to the orthogonal subspaces of the space of flows provided in the Helmholtz decomposition.

Analyzing the projection of a game to the space of potential games may provide useful insights for the original game; see Section 2.7 for a description of ongoing and future work on this direction. We conclude this section by relating the approximate equilibria of a game to the equilibria of the closest potential game.

Theorem 2.6.3. *Let \mathcal{G} be a game, and $\hat{\mathcal{G}}$ be its closest potential game. Assume that h_m denotes the number of strategies of player m , and define $\alpha \triangleq \|\mathcal{G} - \hat{\mathcal{G}}\|_{\mathcal{M},E}$. Then, every ϵ_1 -equilibrium of $\hat{\mathcal{G}}$ is an ϵ -equilibrium of \mathcal{G} for some $\epsilon \leq \max_m \frac{2\alpha}{\sqrt{h_m}} + \epsilon_1$ (and viceversa).*

Proof. By the definition of the norm, it follows that $|u^k(\mathbf{p}) - \hat{u}^k(\mathbf{p})| \leq \frac{1}{\sqrt{h_k}} \|\mathcal{G} - \hat{\mathcal{G}}\|_{\mathcal{M},E} \leq \max_m \frac{\alpha}{\sqrt{h_m}}$, for all $k \in \mathcal{M}$, $\mathbf{p} \in E$. Using Lemma 2.2.1, the result follows. \square

This result implies that the study and characterization of the structure of approximate equilibria in an arbitrary game can be facilitated by making use of the connection between its ϵ -equilibrium set and the equilibria of its closest potential game.

2.7 Summary

We conclude this chapter with a summary of its main contributions. In this chapter, we introduced a novel and natural direct sum decomposition of the space of games into potential, harmonic and nonstrategic subspaces. We studied the equilibrium properties of the subclasses of games induced by this decomposition, and showed that the potential and harmonic components of games have quite distinct and appealing equilibrium properties. In particular, there is a sharp contrast between potential games, that always have pure Nash equilibria, and harmonic games, that generically never do. Moreover, we have shown that while the nonstrategic component does not affect the equilibrium set of games, it can drastically affect their efficiency properties. Using the decomposition framework, we obtained closed-form expressions for the projections of games to their corresponding components, enabling the approximation of arbitrary games by potential and harmonic games. We established that this approximation allows for a systematic method for characterizing the set of ϵ -equilibria of a given game, by relating it to the equilibria of the closest potential game.

2.8 Appendix: Additional Proofs

In this section we provide proofs to some of the results from Sections 2.4 and 2.5.

Proof of Lemmas 2.4.1 and 2.4.2. The proof relies on the fact that $D_m^* D_m = \Delta_{0,m}$ is a Laplacian operator defined on the graph of m -comparable strategy profiles. We show that the kernels of D_m and $\Delta_{0,m}$ coincide, and using the spectral properties of the Laplacian and projection matrices we obtain the desired result.

For a fixed m , it can be seen that strategy profile $\mathbf{p} = (\mathbf{p}^m, \mathbf{p}^{-m})$ is comparable to strategy profiles $(\mathbf{q}^m, \mathbf{p}^{-m})$ for all $\mathbf{q}^m \in E^m$, $\mathbf{q}^m \neq \mathbf{p}^m$ but to none of the strategy profiles $(\mathbf{q}^m, \mathbf{q}^{-m})$ for $\mathbf{q}^{-m} \neq \mathbf{p}^{-m}$. This implies that the graph over which $\Delta_{0,m}$ is defined has $|E^{-m}| = \prod_{k \neq m} h_k$ components (each $\mathbf{p}^{-m} \in E^{-m}$ creates a different component), each of which has $|E^m| = h_m$ elements. Note that all strategy profiles in a component are m -comparable, thus the underlying graph consists of $|E^{-m}|$ components, each of which is a complete graph with $|E^m|$ nodes.

The Laplacian of an unweighted complete graph with n nodes has eigenvalues 0 and n , where the multiplicity of nonzero eigenvalues is $n - 1$ (Chung, 1997). Each component of $\Delta_{0,m}$ leads to eigenvalues 0 and h_m with multiplicities 1 and $h_m - 1$ respectively. Therefore, $\Delta_{0,m}$ has eigenvalues 0 and h_m where the multiplicity of nonzero eigenvalues is $(h_m - 1) \prod_{k \neq m} h_k = \prod_{k \in \mathcal{M}} h_k - \prod_{k \neq m} h_k$. This suggests that the dimension of the kernel of $\Delta_{0,m}$ is $\prod_{k \neq m} h_k$.

Observe that the kernel of $\Delta_{0,m} = D_m^* D_m$ contains the kernel of D_m . For every $\mathbf{q}^{-m} \in E^{-m}$ define $\nu_{\mathbf{q}^{-m}} \in C_0$ such that

$$\nu_{\mathbf{q}^{-m}}(\mathbf{p}) = \begin{cases} 1 & \text{if } \mathbf{p}^{-m} = \mathbf{q}^{-m} \\ 0 & \text{otherwise} \end{cases} \quad (2.48)$$

It is easy to see that $\nu_{\mathbf{p}^{-m}} \perp \nu_{\mathbf{q}^{-m}}$ for $\mathbf{p}^{-m} \neq \mathbf{q}^{-m}$ and $D_m \nu_{\mathbf{p}^{-m}} = 0$ for all $\mathbf{p}^{-m} \in E^{-m}$. Thus, for all \mathbf{q}^{-m} , $\nu_{\mathbf{q}^{-m}}$ belongs to the kernel of D_m and by mutual orthogonality of these functions, the kernel of D_m has dimension at least $|E^{-m}| = \prod_{k \neq m} h_k$. As the dimension of the kernel of $\Delta_{0,m}$ is $\prod_{k \neq m} h_k$ and it contains kernel of D_m , this implies that the kernels of D_m and $\Delta_{0,m}$ coincide.

Thus $\Delta_{0,m}$ maps any $\nu \in C_0$ in the kernel of D_m to zero and scales the ν in the orthogonal complement of the kernel by h_m . On the other hand $D_m^\dagger D_m$ is a projection operator and it has eigenvalue 0 for all functions in the kernel of D_m and 1 for the functions in the orthogonal complement of kernel of D_m . This implies that

$$\Delta_{0,m} = h_m D_m^\dagger D_m, \quad (2.49)$$

and the kernels of $\Delta_{0,m}$, D_m and $D_m^\dagger D_m$ coincide as the claim suggests. \square

Proof of Lemma 2.4.3. For a game, the graph of comparable strategy profiles is connected as can be seen from the definition of the comparable strategy profiles. It is known that for a connected graph, the Laplacian operator has multiplicity 1 for eigenvalue 0 (Chung,

1997). By (2.16) it follows that the function $f \in C_0$ satisfying $f(\mathbf{p}) = 1$ for all $\mathbf{p} \in E$, is an eigenfunction of Δ_0 with eigenvalue 0, implying the result. \square

Proof of Lemma 2.4.4. For the proof of this lemma, we use the following property of the pseudoinverse

$$A^\dagger = (A^*A)^\dagger A^*, \quad (2.50)$$

and the orthogonality properties of the D_m operators: $D_m^*D_k = 0$ and $D_m^\dagger D_k = 0$ if $m \neq k$.

(i) Using (2.50), with $A = D_m$ implies that $D_m^\dagger = (D_m^*D_m)^\dagger D_m^*$. Since for any linear operator L , $(L^\dagger)^* = (L^*)^\dagger$, it follows that $D_m^\dagger = (D_m(D_m^*D_m)^\dagger)^* = (D_m(\Delta_{0,m})^\dagger)^*$. Hence, using Lemma 2.4.1 we obtain $D_m^\dagger = h_m(D_m(\Pi_m)^\dagger)^*$. Since Π_m is a projection operator to the orthogonal complement of the kernel of D_m , we have $\Pi_m^\dagger = \Pi_m$, and $D_m\Pi_m = D_m$. Hence, it follows that $D_m^\dagger = h_m(D_m\Pi_m)^* = h_m D_m^*$ as claimed.

(ii) The identity in (2.50), implies that

$$\left(\sum_{i \in \mathcal{M}} D_i\right)^\dagger = \left(\left(\sum_{i \in \mathcal{M}} D_i\right)^* \left(\sum_{i \in \mathcal{M}} D_i\right)\right)^\dagger \left(\sum_{i \in \mathcal{M}} D_i\right)^*.$$

By the orthogonality of the image spaces of D_i , it follows that $(\sum_{i \in \mathcal{M}} D_i)^* (\sum_{i \in \mathcal{M}} D_i) = \sum_{i \in \mathcal{M}} D_i^* D_i$, and hence

$$\left(\sum_{i \in \mathcal{M}} D_i\right)^\dagger = \left(\sum_{i \in \mathcal{M}} D_i^* D_i\right)^\dagger \left(\sum_{i \in \mathcal{M}} D_i\right)^*.$$

Right-multiplying the above equation by D_j and using the orthogonality of the image spaces of D_i s it follows that

$$\left(\sum_{i \in \mathcal{M}} D_i\right)^\dagger D_j = \left(\sum_{i \in \mathcal{M}} D_i^* D_i\right)^\dagger \left(\sum_{i \in \mathcal{M}} D_i\right)^* D_j = \left(\sum_{i \in \mathcal{M}} D_i^* D_i\right)^\dagger D_j^* D_j.$$

(iii) From the definition of pseudoinverse, it is sufficient to show the following 4 properties to prove the claim: a) $(DD^\dagger)^* = DD^\dagger$, b) $(D^\dagger D)^* = D^\dagger D$, c) $DD^\dagger D = D$, d) $D^\dagger DD^\dagger = D^\dagger$.

Using the identity $D_m^\dagger D_k = 0$ for $k \neq m$, it follows that $DD^\dagger = \sum_{m \in \mathcal{M}} D_m D_m^\dagger$, and $D^\dagger D = \text{diag}(D_1^\dagger D_1, \dots, D_M^\dagger D_M)$. The pseudoinverse of D_m satisfies the properties $D_m^\dagger D_m = (D_m^\dagger D_m)^*$ and $D_m D_m^\dagger = (D_m D_m^\dagger)^*$, and the requirements a) and b) follow immediately using these properties. The identity $D_m^\dagger D_k = 0$ also implies that $DD^\dagger D = [D_1 D_1^\dagger D_1, \dots, D_M D_M^\dagger D_M]$, and $D^\dagger DD^\dagger = [D_1^\dagger D_1 D_1^\dagger; \dots; D_M^\dagger D_M D_M^\dagger]$. Since the pseudoinverse of D_m also satisfies $D_m^\dagger D_m D_m^\dagger = D_m^\dagger$, and $D_m D_m^\dagger D_m = D_m$, the requirements c) and d) are satisfied and the claim follows.

(iv) Since $\Pi = \text{diag}(\Pi_1, \dots, \Pi_M)$, and $\Pi_m = D_m^\dagger D_m$, it follows that

$$D^\dagger D = \text{diag}(D_1^\dagger D_1, \dots, D_M^\dagger D_M) = \text{diag}(\Pi_1, \dots, \Pi_M) = \Pi.$$

(v) Using the identities $D_m^\dagger D_k = 0$ for $k \neq m$, $\delta_0 = \sum_{m \in \mathcal{M}} D_m$, it follows that

$$DD^\dagger \delta_0 = DD^\dagger \sum_{m \in \mathcal{M}} D_m = \sum_{m \in \mathcal{M}} D_m D_m^\dagger D_m = \sum_{m \in \mathcal{M}} D_m = \delta_0.$$

□

Proof of Proposition 2.4.1. Lemma 2.4.2 provides a basis for kernel of D_m and $\dim(\ker(D_m)) = |E^{-m}|$, i.e., the cardinality of the basis is equal to $|E^{-m}|$. By definition $\mathcal{N} = \ker D = \prod_{m \in \mathcal{M}} \ker(D_m)$, hence

$$\dim(\mathcal{N}) = \sum_{m \in \mathcal{M}} \dim(\ker(D_m)) = \sum_{m \in \mathcal{M}} |E^{-m}| = \sum_{m \in \mathcal{M}} \prod_{k \neq m} h_k. \quad (2.51)$$

Next consider the subspace \mathcal{P} of normalized potential games. By definition, the games in this set generate globally consistent flows. Moreover, by Lemma 2.4.6 it follows that there is a unique game in \mathcal{P} , which generates a given gradient flow. Thirdly, note that any globally consistent flow can be obtained as $\delta_0 \phi$ for some $\phi \in C_0$, and the game $\{\Pi_m \phi\}_{m \in \mathcal{M}} \in \mathcal{P}$ generates the same flows as $\delta_0 \phi$. These three facts imply that there is a linear bijective mapping between the games in \mathcal{P} and the globally consistent flows, and hence the dimension of \mathcal{P} is equal to the dimension of the globally consistent flows.

On the other hand, the dimension of the globally consistent flows is equivalent to $\dim(\text{im}(\delta_0))$. Since $\Delta_0 = \delta_0^* \delta_0$ it follows that $\ker(\delta_0) \subset \ker(\Delta_0)$. By Lemma 2.4.3 it follows that $\ker(\Delta_0) = \{f \in C_0 \mid f(\mathbf{p}) = c \in \mathbb{R}, \text{ for all } \mathbf{p} \in E\}$. It follows from the definition of δ_0 that $\delta_0 f = 0$ for all $f \in \ker(\Delta_0)$. These facts imply that $\ker(\delta_0) = \ker(\Delta_0)$ and hence $\dim(\ker(\delta_0)) = 1$. Since δ_0 is a linear operator it follows that $\dim(\text{im}(\delta_0)) = \dim(C_0) - \dim(\ker(\delta_0)) = |E| - 1 = \prod_{m \in \mathcal{M}} h_m - 1$.

Finally observe that $\dim(\mathcal{G}_{\mathcal{M}, E}) = \dim(C_0^M) = M|E| = M \prod_{m \in \mathcal{M}} h_m$. Theorem 2.4.1 implies that $\dim(\mathcal{G}_{\mathcal{M}, E}) = \dim(\mathcal{P}) + \dim(\mathcal{H}) + \dim(\mathcal{N})$. Therefore, it follows that $\dim(\mathcal{H}) = (M - 1) \prod_{m \in \mathcal{M}} h_m - \sum_{m \in \mathcal{M}} \prod_{k \neq m} h_k + 1$. □

Proof of Lemma 2.5.2. Let $X = Du$ denote the pairwise comparison function of the harmonic game. By definition, $(\delta_0^* X)(\mathbf{p}) = 0$ for all $\mathbf{p} \in E$. Thus, for all $\mathbf{r}^m \in E^m$, it follows that

$$0 = \sum_{\mathbf{p}^{-m} \in E^{-m}} (\delta_0^* X)(\mathbf{r}^m, \mathbf{p}^{-m}) = \sum_{\mathbf{p} \in S} (\delta_0^* X)(\mathbf{p}) \quad (2.52)$$

where $S = \{(\mathbf{r}^m, \mathbf{p}^{-m}) \mid \mathbf{p}^{-m} \in E^{-m}\}$. To complete the proof we require the following identity related to the pairwise comparison functions.

Lemma 2.8.1. For all $\hat{X} \in C_1$ and set of strategy profiles $\hat{S} \subset E$, $\sum_{\mathbf{p} \in \hat{S}} (\delta_0^* \hat{X})(\mathbf{p}) = -\sum_{\mathbf{p} \in \hat{S}} \sum_{\mathbf{q} \in \hat{S}^c} \hat{X}(\mathbf{p}, \mathbf{q})$.

Proof. It follows from the definition of δ_0^* that

$$\sum_{\mathbf{p} \in \hat{S}} (\delta_0^* \hat{X})(\mathbf{p}) = -\sum_{\mathbf{p} \in \hat{S}} \sum_{\mathbf{q} \in E} \hat{X}(\mathbf{p}, \mathbf{q}) = -\sum_{\mathbf{p} \in \hat{S}} \sum_{\mathbf{q} \in \hat{S}^c} \hat{X}(\mathbf{p}, \mathbf{q}) - \sum_{\mathbf{p} \in \hat{S}} \sum_{\mathbf{q} \in \hat{S}} \hat{X}(\mathbf{p}, \mathbf{q}) = -\sum_{\mathbf{p} \in \hat{S}} \sum_{\mathbf{q} \in \hat{S}^c} \hat{X}(\mathbf{p}, \mathbf{q}).$$

since $\hat{X}(\mathbf{p}, \mathbf{q}) + \hat{X}(\mathbf{q}, \mathbf{p}) = 0$ for any \mathbf{p}, \mathbf{q} and thus $\sum_{\mathbf{p} \in \hat{S}} \sum_{\mathbf{q} \in \hat{S}} \hat{X}(\mathbf{p}, \mathbf{q}) = 0$. \square

Using this lemma in (2.52), we obtain

$$0 = -\sum_{\mathbf{p} \in S} \sum_{\hat{\mathbf{p}} \in S^c} X(\mathbf{p}, \hat{\mathbf{p}}) = \sum_{\mathbf{p}^{-m} \in E^{-m}} \sum_{\mathbf{p}^m \in E^m} u^m(\mathbf{r}^m, \mathbf{p}^{-m}) - u^m(\mathbf{p}^m, \mathbf{p}^{-m}). \quad (2.53)$$

Since \mathbf{r}^m is arbitrary, it follows that $\sum_{\mathbf{p}^{-m} \in E^{-m}} u^m(\mathbf{q}^m, \mathbf{p}^{-m}) - u^m(\mathbf{r}^m, \mathbf{p}^{-m}) = 0$ for all $\mathbf{q}^m, \mathbf{r}^m \in E^m$. \square

Proof of Theorem 2.5.8. Since $\mathcal{Z} \oplus \mathcal{I} = \mathcal{G}_{\mathcal{M}, E}$, the last column immediately follows from Proposition 2.5.1 and Theorem 2.4.1. Below, we present the dimension results for each row of the table, and the corresponding entries in the first two columns.

Throughout the proof we denote by e the h dimensional vector of ones. Since $\mathcal{N} = \ker D$, the two-player games in \mathcal{N} take the form (ea^T, be^T) for some $a, b \in \mathbb{R}^h$. We shall make use of this fact in the proof.

$\mathcal{P} \oplus \mathcal{H} \oplus \mathcal{N}$: Since $\mathcal{P} \oplus \mathcal{H} \oplus \mathcal{N} = \mathcal{G}_{\mathcal{M}, E}$, it follows that $\dim(\mathcal{P} \oplus \mathcal{H} \oplus \mathcal{N}) \cap \mathcal{Z} = \dim \mathcal{Z}$ and $\dim(\mathcal{P} \oplus \mathcal{H} \oplus \mathcal{N}) \cap \mathcal{I} = \dim \mathcal{I}$. Zero-sum games are games with payoff matrices $(A, -A)$ for some $A \in \mathbb{R}^{h \times h}$. Thus, the dimension of the zero-sum games is equivalent to the dimension of possible A matrices that define zero-sum games and hence $\dim \mathcal{Z} = h^2$. Similarly, identical interest games are games with payoff matrices (A, A) for some $A \in \mathbb{R}^{h \times h}$, and hence $\dim \mathcal{I} = h^2$.

$\mathcal{P} \oplus \mathcal{N}$: By Theorem 2.5.1, it follows that $\mathcal{P} \oplus \mathcal{N}$ is equivalent to the set of potential games. Observe that all identical interest games are potential games, where the utility functions of players are equal to the potential function of the game. Thus, it follows that $\dim(\mathcal{P} \oplus \mathcal{N}) \cap \mathcal{I} = \dim \mathcal{I} = h^2$.

Let \mathcal{G} denote a zero-sum game in $\mathcal{P} \oplus \mathcal{N}$, with payoff matrices $(A, -A)$, and denote the matrix corresponding to a potential function of \mathcal{G} by ϕ . Thus, both the game with payoffs $(A, -A)$ and (ϕ, ϕ) belong to $\mathcal{N} \oplus \mathcal{P}$, and $(A, -A)$ is different from (ϕ, ϕ) by its nonstrategic component. Hence, for some $a, b \in \mathbb{R}^h$, $A = \phi + ea^T$, $-A = \phi + be^T$, for some $a, b \in \mathbb{R}^h$ and

$$A - A = \phi + ea^T + \phi + be^T = 2\phi + ea^T + be^T = 0,$$

thus $-2\phi_{ij} = a_j + b_i$ and $A_{ij} = \phi_{ij} + a_j = \frac{a_j - b_i}{2}$ for all $i, j \in \{1 \dots n\}$. Hence, $a, b \in \mathbb{R}^h$ characterize the possible payoff matrices A , and it can be seen that the set of these matrices

has dimension $2h - 1$. Since these matrices uniquely characterize zero-sum games that are also potential games, it follows that the dimension of $(\mathcal{P} \oplus \mathcal{N}) \cap \mathcal{Z}$ is equal to $2h - 1$.

$\mathcal{H} \oplus \mathcal{N}$: The games in this set do not have potential components. If a game in $\mathcal{H} \oplus \mathcal{N}$ is an identical interest game, then it also belongs to $\mathcal{P} \oplus \mathcal{N}$. Due to the direct sum property of $\mathcal{P} \oplus \mathcal{H} \oplus \mathcal{N}$, it follows that this game can only have nonstrategic component. Therefore, $\dim(\mathcal{H} \oplus \mathcal{N}) \cap \mathcal{I} = \dim \mathcal{N} \cap \mathcal{I}$. Let \mathcal{G} denote a game in $\mathcal{N} \cap \mathcal{I}$. Since \mathcal{G} has only nonstrategic information it follows that its payoffs are given by (ea^T, be^T) , for some vectors a and b . Then, being an identical interest game implies that $ea^T = be^T$, which requires that all entries of payoff matrices are identical, thus $\dim \mathcal{N} \cap \mathcal{I} = 1$.

Consider a zero-sum game in $\mathcal{H} \oplus \mathcal{N}$, with payoff matrices $(A, -A)$. Since, both players have equal number of strategies, the harmonic component of this game is also zero-sum and the payoff matrices in the harmonic component can be denoted by $(A_H, -A_H)$ for some $A_H \in \mathbb{R}^{h \times h}$. Because the original game is in $\mathcal{H} \oplus \mathcal{N}$, the payoff matrices satisfy $A = A_H + ea^T$, $-A = -A_H + be^T$, where (ea^T, be^T) corresponds to the nonstrategic component of the game. It follows that $ea^T + be^T = 0$, and hence ea^T and $-be^T$ are matrices, which have all of their entries identical. Thus, the nonstrategic component of the games in $(\mathcal{H} \oplus \mathcal{N}) \cap \mathcal{Z}$, forms a 1 dimensional subspace. Since the harmonic component is arbitrary, it follows that $\dim(\mathcal{H} \oplus \mathcal{N}) \cap \mathcal{Z} = \dim \mathcal{H} + 1 = (h - 1)^2 + 1 = h^2 - 2h + 2$. \square

Proof of Proposition 2.5.2. We prove the claim, by first showing (i) and (ii) are equivalent and then establishing the equivalence (ii) and (iii).

By the definition of correlated equilibrium, (2.36) implies that x is a correlated equilibrium. To see that any correlated equilibrium of \mathcal{G} satisfies (2.36), assume $x \in \Delta E$ is a correlated equilibrium. Since the game is a harmonic game, by definition, the utility functions $u = \{u^m\}$ satisfy the condition $\delta_0^* Du = 0$. Using (2.13) and (2.21), this condition can equivalently be expressed as

$$\sum_{m \in \mathcal{M}} \sum_{\mathbf{q}^m \in E^m} u^m(\mathbf{q}^m, \mathbf{p}^{-m}) - u^m(\mathbf{p}^m, \mathbf{p}^{-m}) = 0 \quad \text{for all } \mathbf{p} \in E. \quad (2.54)$$

Thus, it follows that

$$\begin{aligned} 0 &= \sum_{\mathbf{p} \in E} x(\mathbf{p}) \sum_{m \in \mathcal{M}} \sum_{\mathbf{q}^m \in E^m} u^m(\mathbf{q}^m, \mathbf{p}^{-m}) - u^m(\mathbf{p}^m, \mathbf{p}^{-m}) \\ &= \sum_{m \in \mathcal{M}} \sum_{\mathbf{q}^m \in E^m} \sum_{\mathbf{p}^m \in E^m} \sum_{\mathbf{p}^{-m} \in E^{-m}} x(\mathbf{p}^m, \mathbf{p}^{-m}) (u^m(\mathbf{q}^m, \mathbf{p}^{-m}) - u^m(\mathbf{p}^m, \mathbf{p}^{-m})). \end{aligned} \quad (2.55)$$

Since x is a correlated equilibrium, $\sum_{\mathbf{p}^{-m} \in E^{-m}} x(\mathbf{p}^m, \mathbf{p}^{-m}) (u^m(\mathbf{q}^m, \mathbf{p}^{-m}) - u^m(\mathbf{p}^m, \mathbf{p}^{-m})) \leq 0$ for all $\mathbf{p}^m, \mathbf{q}^m$ and $m \in \mathcal{M}$. Hence, (2.55) implies that

$$\sum_{\mathbf{p}^{-m} \in E^{-m}} x(\mathbf{p}^m, \mathbf{p}^{-m}) (u^m(\mathbf{q}^m, \mathbf{p}^{-m}) - u^m(\mathbf{p}^m, \mathbf{p}^{-m})) = 0$$

for all $\mathbf{p}^m, \mathbf{q}^m$ and $m \in \mathcal{M}$. Thus, we conclude (i) and (ii) are equivalent.

To see the equivalence of (ii) and (iii), observe that (iii) immediately implies (ii). Assume (ii) holds, then writing (2.36) for two strategies $\mathbf{r}^m, \mathbf{q}^m \in E^m$, and subtracting these equations from each other, it follows that $\sum_{\mathbf{p}^{-m} \in E^{-m}} (u^m(\mathbf{r}^m, \mathbf{p}^{-m}) - u^m(\mathbf{q}^m, \mathbf{p}^{-m})) x(\mathbf{p}^m, \mathbf{p}^{-m}) = 0$. Since \mathbf{r}^m and \mathbf{q}^m are arbitrary it follows that for all $\mathbf{q}^m \in E^m$

$$\sum_{\mathbf{p}^{-m} \in E^{-m}} u^m(\mathbf{q}^m, \mathbf{p}^{-m}) x(\mathbf{p}^m, \mathbf{p}^{-m}) = c_{\mathbf{p}^m}, \quad (2.56)$$

for some $c_{\mathbf{p}^m} \in \mathbb{R}$. Since the game is normalized, we have $\sum_{\mathbf{q}^m \in E^m} u^m(\mathbf{q}^m, \mathbf{p}^{-m}) = 0$. Thus summing (2.56) over $\mathbf{q}^m \in E^m$ it follows that $c_{\mathbf{p}^m} = 0$, and hence (ii) implies (iii).

Therefore we conclude that (i), (ii) and (iii) are equivalent for normalized harmonic games. \square

Note that in the above proof, we used the assumption that the game is normalized, only when establishing the equivalence of (ii) and (iii). Therefore, it can be seen that (i) and (ii) are equivalent for all harmonic games.

2.9 Appendix: Equilibria in Two Player Harmonic Games

In this section we focus on the equilibrium properties of two player harmonic games. We show that sets of correlated and mixed equilibria are identical in these games.

In the following theorem, we present a basis for two-player normalized harmonic games. The idea behind our construction is to obtain a collection of games, in which both players have “effectively” two strategies (the payoffs are equal to zero, if other strategies are played), and ensure that they are linearly independent normalized harmonic games.

Theorem 2.9.1. *Consider the set of two-player games where the first player has h_1 strategies and the second player has h_2 strategies. For any $i \in \{1, \dots, h_1 - 1\}$ and $j \in \{1, \dots, h_2 - 1\}$, define bimatrix games \mathcal{G}^{ij} , with payoff matrices $(h_2 A^{ij}, -h_1 A^{ij})$, where $A^{ij} \in \mathbb{R}^{h_1 \times h_2}$ is such that*

$$A_{kl}^{ij} = \begin{cases} 1 & \text{if } (k, l) = (i, j) \text{ or } (k, l) = (i + 1, j + 1), \\ -1 & \text{if } (k, l) = (i + 1, j) = (k, l) \text{ or } (k, l) = (i, j + 1), \\ 0 & \text{otherwise.} \end{cases} \quad (2.57)$$

The collection $\{\mathcal{G}^{ij}\}$ provides a basis of \mathcal{H} .

Proof. It can be seen that each \mathcal{G}^{ij} is normalized, since row and column sums of A^{ij} is equal to zero. By Theorem 2.5.6 and (2.57), it also follows that \mathcal{G}^{ij} belongs to \mathcal{H} . It can be seen from Proposition 2.4.1 that $\dim \mathcal{H} = (h_1 - 1)(h_2 - 1)$, is equal to the cardinality of the

collection $\{\mathcal{G}^{ij}\}$. Thus, in order to prove the claim, it is sufficient to prove that

$$\sum_{i \in \{1, \dots, h_1-1\}} \sum_{j \in \{1, \dots, h_2-1\}} \alpha_{ij} A^{ij} = 0, \quad (2.58)$$

only if $\alpha_{ij} = 0$ for all i, j .

Note that A^{11} is the only matrix which has a nonzero entry in the first column and the first row. Thus, (2.58) implies that $\alpha_{11} = 0$. Similarly it can be seen that A^{11} and A^{12} are the only matrices which have nonzero entries in the first row and the second column, thus $\alpha_{12} = 0$. Proceeding iteratively it follows that if (2.58) holds, then $\alpha_{ij} = 0$ for all i, j and the claim follows. \square

The next example uses the basis introduced above, to show that in harmonic games, the uniformly mixed strategy profile is not necessarily the unique mixed Nash equilibrium.

Example 2.9.1. *In this example we consider two-player harmonic games, where $E^1 = \{x, y\}$ and $E^2 = \{a, b, c\}$. Using Theorem 2.9.1, a basis for normalized two-player harmonic games is given in Tables 2.7a and 2.7b. Thus, any harmonic game with these strategy sets, can be expressed as in Table 2.7c. Consider some fixed α and β . As can be seen from Definition 2.5.1, the mixed equilibria for this game are given by*

$$\left(\frac{1}{2}, \frac{1}{2}\right) \times (\theta_1, \theta_2, \theta_3)$$

where θ_1, θ_2 and θ_3 are scalars that satisfy $\theta_1 + \theta_2 + \theta_3 = 1$, $\theta_1, \theta_2, \theta_3 \geq 0$ and $\theta_1(6\alpha) + \theta_2(-6\alpha + 6\beta) + \theta_3(-6\beta) = 0$. Note that since there are two linear equations in three variables, this system has a continuum of solutions. Moreover, since $(\theta_1, \theta_2, \theta_3) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is a solution, it follows that there is a continuum of solutions for which $\theta_1, \theta_2, \theta_3 \geq 0$.

Since this is true for any α, β , we conclude that all games in \mathcal{H} have uncountably many mixed equilibria. Additionally, since the nonstrategic component does not affect the equilibrium properties of a game it follows that all harmonic games on $E^1 \times E^2$ (all games in $\mathcal{H} \oplus \mathcal{N}$) have uncountably many mixed Nash equilibria.

	<i>a</i>	<i>b</i>	<i>c</i>
<i>x</i>	3, -2	-3, 2	0, 0
<i>y</i>	-3, 2	3, -2	0, 0

(a) Basis element 1

	<i>a</i>	<i>b</i>	<i>c</i>
<i>x</i>	0, 0	3, -2	-3, 2
<i>y</i>	0, 0	-3, 2	3, -2

(b) Basis element 2

	<i>a</i>	<i>b</i>	<i>c</i>
<i>x</i>	$3\alpha, -2\alpha$	$-3\alpha + 3\beta, 2\alpha - 2\beta$	$-3\beta, 2\beta$
<i>y</i>	$-3\alpha, 2\alpha$	$3\alpha - 3\beta, -2\alpha + 2\beta$	$3\beta, -2\beta$

(c) A game in \mathcal{H}

Table 2.7: Basis of \mathcal{H}

Using this basis, we characterize in the following theorem, the correlated equilibria in two-player harmonic games. Interestingly, our results suggest that in two-player harmonic games, the set of mixed Nash equilibria and correlated equilibria generically coincide.

Theorem 2.9.2. *Consider the set of two-player harmonic games where the first player has h_1 strategies and the second player has h_2 strategies. Without loss of generality assume that $h_1 \geq h_2$. Generically,*

1. *Every correlated equilibrium is a mixed Nash equilibrium, where the player with minimum number of strategies uses the uniformly mixed strategy.*
2. *The dimension of the set of correlated equilibria is $h_1 - h_2$*

Proof. As discussed earlier, nonstrategic components of games do not affect the equilibrium sets. Thus, to prove that (i) and (ii) are generically true for harmonic games, it is sufficient to prove that they generically hold for normalized harmonic games.

Consider a two-player normalized harmonic game with payoff matrices (A, B) , where $A, B \in \mathbb{R}^{h_1 \times h_2}$. By Theorem 2.5.6, it follows that $A = -\frac{h_2}{h_1}B$. Denote by e_1 (similarly e_2), the h_1 (similarly h_2) dimensional vector, all entries of which are identically equal to 1. Since the game is normalized, it follows that $e_1^T A = 0$ and $B e_2 = -\frac{h_1}{h_2} A e_2 = 0$.

Let x be a correlated equilibrium of this game. For each $\mathbf{p}^1 \in E^1$, denote by $x(\mathbf{p}^1, \cdot) \in \mathbb{R}^{h_2}$ the vector of probabilities $[x(\mathbf{p}^1, \mathbf{p}^2)]_{\mathbf{p}^2}$. By Proposition 2.5.2 (iii), it follows that these vectors satisfy the condition

$$Ax(\mathbf{p}^1, \cdot) = 0. \tag{2.59}$$

Note that we need to characterize the kernel of the payoff matrix A , to identify the correlated equilibria. For that reason, we state the following technical lemma:

Lemma 2.9.1. *Consider the set of normalized harmonic games in Theorem 2.9.2. Generically, the payoff matrices of players have their row and column ranks equal to $h_2 - 1$.*

Proof. The payoff matrices of the players satisfy $A = -\frac{h_2}{h_1}B$, so they have the same row and column rank. It follows from Theorem 2.9.1 that the collection of matrices $\{A^{ij}\}$ span the payoff matrices of harmonic games. It can be seen that the matrices in the span of this collection generically have row and column rank equal to $h_2 - 1$, and the claim follows. \square

Using this lemma, it follows that generically the kernel of A is 1 dimensional. As shown earlier, e_2 is in the kernel of A , thus, (2.59), implies that generically $x(\mathbf{p}^1, \cdot)$ has the form $x(\mathbf{p}^1, \cdot) = c_{\mathbf{p}^1} e_2$, for some $c_{\mathbf{p}^1} \in \mathbb{R}$. Since x is a probability distribution, the definition of $x(\mathbf{p}^1, \cdot)$ implies that $x(\mathbf{p}^1, \mathbf{p}^2) = c_{\mathbf{p}^1} \geq 0$, and $\sum_{\mathbf{p}^1 \in E^1, \mathbf{p}^2 \in E^2} x(\mathbf{p}^1, \mathbf{p}^2) = h_2 \sum_{\mathbf{p}^1 \in E^1} c_{\mathbf{p}^1} = 1$. Thus, it follows that $x(\mathbf{p}^1, \mathbf{p}^2) = c_{\mathbf{p}^1} = \frac{\alpha_{\mathbf{p}^1}}{h_2}$, for some $\alpha_{\mathbf{p}^1} \geq 0$ such that $\sum_{\mathbf{p}^1 \in E^1} \alpha_{\mathbf{p}^1} = 1$. It can be seen from this description that generically, the correlated equilibria are mixed equilibria where the first player uses the probability distribution $x^1 = \alpha \triangleq [\alpha_{\mathbf{p}^1}]_{\mathbf{p}^1} \in \Delta E^1$ and the second player uses the distribution $x^2 = \left[\frac{1}{h_2} \right]_{\mathbf{p}^2}$.

Since the correlated equilibria have this form, it can be seen using Proposition 2.5.2 (iii) for the second player that

$$\sum_{\mathbf{p}^1 \in E^1} u^2(\mathbf{p}^1, \mathbf{q}^2) x(\mathbf{p}^1, \mathbf{p}^2) = \frac{1}{h_2} \sum_{\mathbf{p}^1 \in E^1} u^2(\mathbf{p}^1, \mathbf{q}^2) \alpha_{\mathbf{p}^1} = 0, \quad (2.60)$$

where $\alpha \in \Delta E^1$. The above condition can be restated using the payoff matrices as follows:

$$\alpha^T B = -\frac{h_1}{h_2} \alpha^T A = 0, \quad (2.61)$$

where $\alpha \in \Delta E^1$. Since, the row rank of A is $h_2 - 1$, the dimension of α that satisfies (2.61) is $h_1 - h_2 + 1$. Note that since α is a probability distribution, it also satisfies the condition $\alpha^T e_1 = 1$. Note that since $e_1^T A = 0$, this condition is orthogonal to the ones in (2.61). Hence, it follows that the dimension of α which satisfies the correlated equilibrium conditions in (2.61) (other than the positivity) is $h_1 - h_2$. On the other hand, $\alpha = \frac{1}{h_1} e_1$ gives a correlated equilibrium (by Theorem 2.5.5), thus the positivity condition does not change the dimension of the set of correlated equilibria, and the dimension is generically $h_1 - h_2$. \square

An immediate implication of this theorem is the following:

Corollary 2.9.1. *In two-player harmonic games where players have equal number of strategies, the profile of uniformly mixed strategies is generically the unique correlated equilibrium.*

Note that Theorem 2.9.2 implies that in two-player harmonic games, generically there are no correlated equilibria that are not mixed equilibria. This statement fails, when the number of players is more than two, as shown in the following theorem.

Theorem 2.9.3. *Consider a M -player harmonic game, where $M > 2$, and in which each player has h strategies such that $h^M > M(h^2 - 1) + 1$. The set of correlated equilibria is strictly larger than the set of mixed Nash equilibria: The set of correlated equilibria has dimension at least $h^M - 1 - Mh(h - 1)$, and the set of mixed equilibria has dimension at most $M(h - 1)$.*

Proof. Since each player has h strategies, the set of mixed strategies has dimension $M(h - 1)$, and this is a trivial upper bound on the dimension of the set of mixed equilibria. The set of correlated equilibria, on the other hand, is defined by the equalities in Proposition 2.5.2. Note that there are $Mh(h - 1)$ such equalities and the dimension of ΔE is $h^M - 1$, hence the dimension of the correlated equilibria is at least $h^M - 1 - Mh(h - 1)$ (by ignoring possible dependence of the equalities).

The difference in the dimensions implies that the set of correlated equilibria is strictly larger than the set of mixed equilibria. \square

Note that this theorem can be easily generalized to the case when players have different number of strategies. An interesting problem is to find the exact dimensions of the set of mixed Nash and correlated equilibria when there are more than two players. However, due to complicated dependence relations of the correlated equilibrium conditions in Proposition 2.5.2, we do not pursue this question in this chapter, and leave it as a future problem.

Chapter 3

Dynamics in Near-Potential Games

3.1 Introduction

The study of multi-agent strategic interactions both in economics and engineering mainly relies on the concept of Nash equilibrium. This raises the question whether Nash equilibrium makes approximately accurate predictions of the agent behavior. One possible justification for Nash equilibrium is that it arises as the long run outcome of dynamical processes, in which less than fully rational players search for optimality over time. However, unless the game belongs to special (but restrictive) classes of games, such dynamics do not converge to a Nash equilibrium, and there is no systematic analysis of their limiting behavior (Jordan, 1993; Fudenberg and Levine, 1998; Shapley, 1964).

In potential games, which we discussed in the previous chapter, many of the simple user dynamics, such as best response dynamics and fictitious play, converge to a Nash equilibrium (Monderer and Shapley, 1996b,a; Fudenberg and Levine, 1998; Sandholm, 2010a; Young, 2004). Intuitively, dynamics in potential games and dynamics in games that are “close” (in terms of the payoffs of the players) to potential games should be related. Our goal in this chapter is to make this intuition precise and provide a systematic framework for studying dynamics in finite strategic form games by exploiting their relation to close potential games.

We start by illustrating via examples that general games which are close in terms of payoffs may have significantly different limiting behavior under simple user dynamics. Our first example focuses on better response dynamics in which at each step or strategy profile, a player (chosen consecutively or at random) updates its strategy unilaterally to one that yields a better payoff.¹

¹Consider a game where players are not indifferent between their strategies at any strategy profile. Arbitrarily small payoff perturbations of this game lead to games which have the same better response structure as the original game. Hence, for a given game there may exist a close enough game such that the outcome of the better response dynamics in two games are identical. However, for payoff differences of given size it is always possible to find games with different better response properties as illustrated in Example 3.1.1.

Example 3.1.1. Consider two games with two players and payoffs given in Figure 3-1. The entries of these tables indexed by row X and column Y show payoffs of the players when the first player uses strategy X and the second player uses strategy Y . Let $0 < \theta \ll 1$. Both games have a unique Nash equilibrium: (B, B) for \mathcal{G}_1 , and the mixed strategy profile $\left(\frac{2}{3}A + \frac{1}{3}B, \frac{\theta}{1+\theta}A + \frac{1}{1+\theta}B\right)$ for \mathcal{G}_2 .

We consider convergence of the sequence of pure strategy profiles generated by the better response dynamics. In \mathcal{G}_1 , the sequence converges to strategy profile (B, B) . In \mathcal{G}_2 , the sequence does not converge (it can be shown that the sequence follows the better response cycle (A, A) , (B, A) , (B, B) and (A, B)). Thus, trajectories are not contained in any ϵ -equilibrium set for $\epsilon < 2$.

	A	B
A	0, 1	0, 0
B	1, 0	$\theta, 2$

\mathcal{G}_1

	A	B
A	0, 1	0, 0
B	1, 0	$-\theta, 2$

\mathcal{G}_2

Figure 3-1: A small change in payoffs results in significantly different behavior for the pure strategy profiles generated by the better response dynamics.

The second example considers fictitious play dynamics, where at each step, each player maintains an (independent) empirical frequency distribution of other player's strategies and plays a best response against it.

Example 3.1.2. Consider two games with two players and payoffs given in Figure 3-2. Let θ be an irrational number such that $0 < \theta \ll 1$. It can be seen that \mathcal{G}_1 has multiple equilibria (including pure equilibria (A, A) , (B, B) and (C, C)), whereas \mathcal{G}_2 has a unique equilibrium given by the mixed strategy profile where both players assign $1/3$ probability to each of its strategies.

	A	B	C
A	$1 + \theta, 1 + \theta$	1, 0	0, 1
B	0, 1	$1 + \theta, 1 + \theta$	1, 0
C	1, 0	0, 1	$1 + \theta, 1 + \theta$

\mathcal{G}_1

	A	B	C
A	$1 - \theta, 1 - \theta$	1, 0	0, 1
B	0, 1	$1 - \theta, 1 - \theta$	1, 0
C	1, 0	0, 1	$1 - \theta, 1 - \theta$

\mathcal{G}_2

Figure 3-2: A small change in payoffs results in significantly different behavior for the empirical frequencies generated by the fictitious play dynamics.

We focus on the convergence of the sequence of empirical frequencies generated by the fictitious play dynamics (under the assumption that initial empirical frequency distribution assigns probability 1 to a pure strategy profile). In \mathcal{G}_1 , this sequence converges to a pure equilibrium starting from any pure strategy profile. In \mathcal{G}_2 , the sequence displays oscillations similar to those seen in the Shapley game (see Shapley (1964); Fudenberg and Levine (1998)). To see this, assume that the initial empirical frequency distribution assigns probability 1 to the strategy profile (A, A) . Observe that since the underlying game is a symmetric game, empirical frequency distribution of each player will be identical at all steps. Starting from (A, A) , both players update their strategy to C . After sufficiently many updates, the empirical frequency of A falls below $\theta/(1 + \theta)$, and that of C exceeds $1/(1 + \theta)$. Thus, the payoff specifications suggest that both players start using strategy B . Similarly, after empirical frequency of B exceeds $1/(1 + \theta)$, and that of C falls below $\theta/(1 + \theta)$, then both players start playing A . Observe that update to a new strategy takes place only when one of the strategies is being used with very high probability (recall that $\theta \ll 1$) and this feature of empirical frequencies is preserved throughout. For this reason the sequence of empirical frequencies does not converge to $(1/3, 1/3, 1/3)$, the unique Nash equilibrium of \mathcal{G}_2 .

These examples suggest that in general, it may not be possible to characterize the limiting dynamics in a given game, by using knowledge of the limiting behavior in a nearby game. In this chapter, in contrast with this observation, we will show that games that are close (in terms of payoffs of players) to potential games have similar limiting dynamics to those in potential games. Moreover, it is possible to provide a *quantitative measure of the size of the limiting set of dynamics* in terms of the ‘distance’ of the game from potential games. Our approach relies on using the potential function of a close potential game for the analysis of commonly studied update rules.² We note that our results hold for arbitrary strategic form games, however our characterization of limiting behavior of dynamics is more informative for games that are close to potential games. We therefore focus our investigation to such games in this chapter and refer to them as *near-potential games*.

We start our analysis by introducing *maximum pairwise difference*, a measure of “closeness” of games.³ Let \mathbf{p} and \mathbf{q} be two strategy profiles, which differ in the strategy of a single player, say player m . Recall that we refer to the change in the payoff of player m between these two strategy profiles, as the pairwise comparison of \mathbf{p} and \mathbf{q} (see Section 2.2.2). Intuitively, this quantity captures how much player m can improve its utility by unilaterally deviating from strategy profile \mathbf{p} to strategy profile \mathbf{q} . For given games, the maximum pairwise difference is defined as the maximum difference between the pairwise comparisons of these games. Thus, the maximum pairwise difference captures how different two games are in terms of the utility improvements due to unilateral deviations. Since equi-

²Throughout the chapter, we use the terms *learning dynamics* and *update rules* interchangeably.

³Maximum pairwise difference is closely related to the norm we introduced for space of games in Section 2.6. However, it is preferred in this section, since it allows us to focus on the “strategic” difference between games, as we explain in detail in Section 3.2.

libria of games, and strategy updates in various update rules (such as better/best response dynamics) can be expressed in terms of unilateral deviations, maximum pairwise difference provides a measure of strategic similarities of games. As established in Chapter 2 the set of potential games is a subspace of space of games. Using this observation, we show that the closest potential game to a given game, in the sense of maximum pairwise difference, can be obtained by solving a convex optimization problem. This provides an alternative way (that is different than the decomposition approach) of approximating a given game with a potential game that has a similar equilibrium set and dynamic properties, as illustrated in Example 3.1.3.

Example 3.1.3. Consider a two-player game \mathcal{G} , which is not a potential game, and the closest potential game to this game (in terms of maximum pairwise difference), $\hat{\mathcal{G}}$, given in Figure 3-3. The maximum pairwise difference of these games is 2, since the utility improvements in these games due to unilateral deviations differ by at most 2 (For instance consider the deviation of the column player from (A, A) to (A, B) . In \mathcal{G} this leads to a utility improvement of 6, whereas, in $\hat{\mathcal{G}}$ the improvement amount is 4). It can be seen that for both games (B, B) is the unique equilibrium. Moreover, trajectories of better response dynamics and empirical frequencies of fictitious play dynamics converge to this equilibrium in both games.

	A	B
A	8, 2	8, 8
B	2, 2	12, 10

\mathcal{G}

	A	B
A	7, 3	9, 7
B	3, 1	11, 11

$\hat{\mathcal{G}}$

Figure 3-3: A game (\mathcal{G}) and a nearby potential game ($\hat{\mathcal{G}}$) share similar equilibrium set and dynamic properties.

We focus on three commonly studied user dynamics: discrete-time better/best response, logit response, and discrete-time fictitious play dynamics, and establish different notions of convergence for each. We first study *better/best response dynamics*. It is known that the sequence of pure strategy profiles, which we refer to as *trajectories*, generated by these update rules converge to pure Nash equilibria in potential games (Monderer and Shapley, 1996b; Young, 2004). In near-potential games, a pure Nash equilibrium need not even exist. For this reason we focus on the notion of *pure approximate equilibria* or ϵ -*equilibria*, and show that in near-potential games trajectories of these update rules converge to a pure approximate equilibrium set. The size of this set only depends on the distance from the original game to a potential game, and is independent of the payoffs in the original game. In particular, our result for better/best response dynamics establish a ‘Lipschitz-type’ property, i.e., we can find a constant h (which is equal to the number of strategy profiles in the game as shown in Theorem 3.3.1) such that in a game that is δ different (in

terms of maximum pairwise difference) from a potential game the trajectory converges to the δh -equilibrium set.

We then focus on *logit response* dynamics. With this update rule, agents, when updating their strategies, choose their best responses with high probability, but also explore other strategies with a nonzero probability. Logit response induces a Markov chain on the set of pure strategy profiles. The stationary distribution of this Markov chain is used to explain the limiting behavior of this update rule (Young, 1993; Blume, 1997, 1993; Alós-Ferrer and Netzer, 2010; Marden and Shamma, 2012). In potential games, the stationary distribution can be expressed in closed form in terms of the potential function of the game. Additionally, the *stochastically stable strategy profiles*, i.e., the strategy profiles which have nonzero stationary distribution as the exploration probability goes to zero, are those that maximize the potential function (Alós-Ferrer and Netzer, 2010; Blume, 1997; Marden and Shamma, 2012). Exploiting their relation to close potential games, we obtain similar results for near-potential games: (i) we obtain an explicit characterization of the stationary distribution in terms of the distance of the game from a close potential game and the corresponding potential function, and (ii) we show that the stochastically stable strategy profiles are the strategy profiles that approximately maximize the potential of a close potential game, implying that they are pure approximate equilibria of the game. Our analysis relies on a novel perturbation result for Markov chains (see Theorem 3.4.1) which provides bounds on deviations from a stationary distribution when transition probabilities of a Markov chain are *multiplicatively* perturbed, and therefore may be of independent interest.

A summary of our convergence results on better/best response and logit response dynamics can be found in Table 3.1.

Update Rule	Convergence Result
Better/Best Response Dynamics	(Theorem 3.3.1) Trajectories of dynamics converge to $\mathcal{X}_{\delta h}$, i.e., the δh -equilibrium set of \mathcal{G} .
Logit Response Dynamics (with parameter τ)	(Corollary 3.4.2) Stationary distribution μ_τ of logit response dynamics is such that $\left \mu_\tau(\mathbf{p}) - \frac{e^{\frac{1}{\tau}\phi(\mathbf{p})}}{\sum_{\mathbf{q} \in E} e^{\frac{1}{\tau}\phi(\mathbf{q})}} \right \leq \frac{e^{\frac{2\delta(h-1)}{\tau}} - 1}{e^{\frac{2\delta(h-1)}{\tau}} + 1}$, for all \mathbf{p} .
Logit Response Dynamics	(Corollary 3.4.3) Stochastically stable strategy profiles of \mathcal{G} are (i) contained in $S = \{\mathbf{p} \phi(\mathbf{p}) \geq \max_{\mathbf{q}} \phi(\mathbf{q}) - 4\delta(h-1)\}$, (ii) $4\delta h$ -equilibria of \mathcal{G} .

Table 3.1: Convergence properties of better/best response and logit response dynamics in near-potential games. Given a game \mathcal{G} , we use $\hat{\mathcal{G}}$ to denote a nearby potential game with potential function ϕ such that the distance (in terms of the maximum pairwise difference, defined in Section 3.2) between the two games is δ . We use the notation \mathcal{X}_ϵ to denote the ϵ -equilibrium set of the original game, h to denote the number of strategy profiles, μ_τ and $\hat{\mu}_\tau$ to denote the stationary distributions of logit response dynamics in \mathcal{G} and $\hat{\mathcal{G}}$, respectively.

We finally analyze *fictitious play dynamics* in near-potential games. In potential games trajectories of fictitious play need not converge to a Nash equilibrium, but the empirical

frequencies of the played strategies converge to a (mixed) Nash equilibrium (Monderer and Shapley, 1996a; Shamma and Arslan, 2004). In our analysis of fictitious play dynamics, we first show that in near-potential games if the empirical frequencies are outside some ϵ -equilibrium set, then the potential of the close potential game (evaluated at the empirical frequency distribution) increases with each strategy update. Using this result we establish convergence of fictitious play dynamics to a set which can be characterized in terms of the ϵ -equilibrium set of the game and the level sets of the potential function of a close potential game. This result suggests that in near-potential games, the empirical frequencies of fictitious play converge to a set of mixed strategies that (in the close potential game) have potential almost as large as the potential of Nash equilibria. Moreover, exploiting the property that for small ϵ , ϵ -equilibria are contained in disjoint neighborhoods of equilibria, we strengthen our result and establish that if a game is sufficiently close to a potential game, then empirical frequencies of fictitious play dynamics converge to a small neighborhood of equilibria, whose size is explicitly characterized.⁴ Our result recovers as a special case convergence of empirical frequencies to Nash equilibria in potential games.⁵

A summary of our results on convergence of fictitious play dynamics is given in Table 3.2.

The framework provided in this chapter enables us to study the limiting behavior of adaptive user dynamics in arbitrary finite strategic form games. In particular, for a given game we can use either the proposed convex optimization formulation, or the decomposition approach introduced in Chapter 2, to find a nearby potential game and use the distance between these games to obtain a quantitative characterization of the limiting approximate equilibrium set. The characterization this approach provides will be tighter if the original game is closer to a potential game.

3.1.1 Related Literature

In potential games, pure Nash equilibria are stable under various learning dynamics such as better/best response dynamics (Monderer and Shapley, 1996b; Fudenberg and Levine, 1998; Young, 2004). Because of these properties, potential games found applications in various control and resource allocation problems (Monderer and Shapley, 1996b; Marden et al., 2009a; Candogan et al., 2010b; Arslan et al., 2007).

⁴The bounds we obtain for the limiting behavior of fictitious play dynamics have a different flavor than those for better/best response dynamics, and logit response. While the bounds we obtain for the latter update rules are independent of the payoffs (and a function of only δ and the number of strategy profiles in the game), for fictitious play they are not. This is because, fictitious play results exploit the structure of mixed (approximate) equilibrium sets, which rely on the actual payoff parameters, whereas other dynamics results do not involve mixed strategies.

⁵This result also implies that in near-potential games fictitious play dynamics are upper semicontinuous with respect to payoff parameters. This upper semicontinuity result could alternatively be proved by considering differential inclusions that represent the limiting behavior of fictitious play (see Benaïm et al. (2005)), together with upper semicontinuity results on differential inclusions (Li and Zhang, 2002). Our result, in addition to upper semicontinuity, provides explicit bounds on the size of the limiting set.

Update Rule	Convergence Result
Fictitious Play	(Corollary 3.5.1) Empirical frequencies of dynamics converge to the set of mixed strategies with large enough potential: $\{\mathbf{x} \in \prod_m \Delta E^m \mid \phi(\mathbf{x}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y})\}$
Fictitious Play	<p>(Theorem 3.5.2) Assume that \mathcal{G} has finitely many equilibria. There exists some $\bar{\delta} > 0$, and $\bar{\epsilon} > 0$ (which are functions of utilities of \mathcal{G} but not δ) such that if $\delta < \bar{\delta}$, then the empirical frequencies of fictitious play converge to</p> $\left\{ \mathbf{x} \mid \ \mathbf{x} - \mathbf{x}_k\ \leq \frac{4f(M\delta)ML}{\epsilon} + f(M\delta + \epsilon), \right.$ <p style="text-align: center;">for some equilibrium \mathbf{x}_k $\left. \right\}$,</p> <p>for any ϵ such that $\bar{\epsilon} \geq \epsilon > 0$, where $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is an upper semicontinuous function that quantifies the size of mixed equilibrium sets (defined explicitly in Section 3.5) such that $f(x) \rightarrow 0$ as $x \rightarrow 0$.</p>

Table 3.2: Convergence properties of fictitious play dynamics in near-potential games. We denote the number of players in the game by M , set of mixed strategies of player m by ΔE^m , and the Lipschitz constant of the mixed extension of ϕ by L . Rest of the notation is the same as in Table 3.1.

There is no systematic framework for analyzing the limiting behavior of many of the adaptive update rules in general games (Jordan, 1993; Fudenberg and Levine, 1998; Shapley, 1964). However, for potential games there is a long line of literature establishing convergence of natural adaptive dynamics such as better/best response dynamics (Monderer and Shapley, 1996b; Young, 2004), fictitious play (Monderer and Shapley, 1996a; Shamma and Arslan, 2004; Marden et al., 2009b; Hofbauer and Sandholm, 2002) and logit response dynamics (Blume, 1993, 1997; Alós-Ferrer and Netzer, 2010; Marden and Shamma, 2012).

Another strand of literature focuses on identifying classes of games with similar properties to potential games. Examples include ordinal potential games (Monderer and Shapley, 1996b), best-response potential games (Voorneveld, 2000), pseudo-potential games (Dubey et al., 2006), and nested potential games (Uno, 2007). Even though these classes of games share similar ordinal properties with potential games, for update processes that involve mixed strategies (such as fictitious play), or that rely on actual payoff values (such as logit response), they do not lead to simple analysis unless further structure is imposed (unlike potential games). For this reason, in this chapter we follow a different approach, and characterize dynamic properties of games, by exploiting their closeness to potential games.

There are also papers in the literature, which identify classes of games that are strategically equivalent to potential games (Morris and Ui, 2004). These equivalence notions can be used to extend the dynamical properties of potential games to their equivalence

classes. However, we want to emphasize that the framework presented in this chapter can be applied for study of dynamics in games that are not strategically equivalent to potential games, thereby providing tools for study of dynamics in arbitrary strategic form games.

3.1.2 Outline

The rest of the chapter is organized as follows. Most of the notation, and game theoretic preliminaries we need for this section are borrowed from Chapter 2. In Section 3.2 we introduce some additional notation, and cover the game theoretic preliminaries that were not discussed in the previous chapter. We present an analysis of better and best response dynamics in near-potential games in Section 3.3. In Section 3.4, we extend our analysis to logit response, and focus on the stationary distribution and stochastically stable states of logit response. We present the results on fictitious play in Section 3.5. We close in Section 3.6 with a summary of the main contributions of this chapter. Some of the proofs are delegated to Section 3.7.

3.2 Preliminaries

In this section, we present the game-theoretic background that is relevant to this chapter. Additionally, we introduce the closeness measure for games, which is used in the rest of the chapter.

Our focus in this chapter is on finite strategic form games. We inherit the notation used in Chapter 2 for such games, and denote them using the tuple $\langle \mathcal{M}, \{E^m\}, \{u^m\} \rangle$. In subsequent sections, we characterize the limiting behavior of dynamics in finite games in relation to their equilibria. In particular, we make use of Nash equilibria, and ϵ -equilibria (see Section 2.2.1 for definitions). We denote the set of ϵ -equilibria in a game \mathcal{G} by \mathcal{X}_ϵ .

In order to establish our results on dynamics, we exploit the properties of potential games (see Section 2.2.1 for a definition). We next explain an important property of potential games, which plays a key role in our analysis: In potential games the total unilateral utility improvement around a “closed path” is equal to zero. Before we formally state this result, we first provide some necessary definitions, which are also used in Section 3.3 when we analyze better/best response dynamics in near-potential games.

Definition 3.2.1 (Path – Closed Path – Improvement Path). *A path is a collection of strategy profiles $\gamma = (\mathbf{p}_0, \dots, \mathbf{p}_N)$ such that \mathbf{p}_i and \mathbf{p}_{i+1} differ in the strategy of exactly one player. A path is a closed path (or a cycle) if $\mathbf{p}_0 = \mathbf{p}_N$. A path is an improvement path if $u^{m_i}(\mathbf{p}_i) \geq u^{m_i}(\mathbf{p}_{i-1})$ where m_i is the player who modifies its strategy when the strategy profile is updated from \mathbf{p}_{i-1} to \mathbf{p}_i .*

The transition from strategy profile \mathbf{p}_{i-1} to \mathbf{p}_i is referred to as *step i of the path*. The length of a path is equal to its number of steps, i.e., the length of the path $\gamma = (\mathbf{p}_0, \dots, \mathbf{p}_N)$

is N . We say that a closed path is *simple* if no strategy profile other than the first and the last strategy profiles is repeated along the path. For any path $\gamma = (\mathbf{p}_0, \dots, \mathbf{p}_N)$ let $I(\gamma)$ represent the total utility improvement along the path, i.e.,

$$I(\gamma) = \sum_{i=1}^N u^{m_i}(\mathbf{p}_i) - u^{m_i}(\mathbf{p}_{i-1}),$$

where m_i is the index of the player that modifies its strategy in the i th step of the path. The following proposition provides a necessary and sufficient condition under which a given game is a potential game.

Proposition 3.2.1 (Monderer and Shapley (1996b)). *A game is a potential game if and only if $I(\gamma) = 0$ for all simple closed paths γ .*

We next provide a formal definition of the measure of “closeness” of games, used in the subsequent sections.

Definition 3.2.2 (Maximum Pairwise Difference). *Let \mathcal{G} and $\hat{\mathcal{G}}$ be two games with set of players \mathcal{M} , set of strategy profiles E , and collections of utility functions $\{u^m\}_{m \in \mathcal{M}}$ and $\{\hat{u}^m\}_{m \in \mathcal{M}}$ respectively. The maximum pairwise difference (MPD) between these games is defined as*

$$d(\mathcal{G}, \hat{\mathcal{G}}) \triangleq \max_{\mathbf{p} \in E, m \in \mathcal{M}, q^m \in E^m} |(u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})) - (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m}))|.$$

Note that the pairwise difference $u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})$ quantifies how much player m can improve its utility by unilaterally deviating from strategy profile (p^m, \mathbf{p}^{-m}) to strategy profile (q^m, \mathbf{p}^{-m}) . Thus, the MPD captures how different two games are in terms of the utility improvements due to unilateral deviations.⁶ We refer to pairs of games with small MPD as *close games*, and games that have a small MPD to a potential game as *near-potential games*.

The MPD measures the closeness of games in terms of the difference of unilateral deviations, rather than the difference of their utility functions, i.e., quantities of the form

$$|(u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})) - (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m}))|$$

⁶ An alternative distance measure can be given by

$$d_2(\mathcal{G}, \hat{\mathcal{G}}) \triangleq \left(\sum_{\mathbf{p} \in E} \sum_{m \in \mathcal{M}, q^m \in E^m} ((u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})) - (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m})))^2 \right)^{\frac{1}{2}},$$

and this quantity corresponds to the 2-norm of the difference of \mathcal{G} and $\hat{\mathcal{G}}$ in terms of the utility improvements due to unilateral deviations. Our analysis of the limiting behavior of dynamics relies on the maximum of such utility improvement differences between a game and a near-potential game. Thus, the measure in Definition 3.2.2 provides tighter bounds for our dynamics results, and hence is preferred in this chapter.

are used to identify close games, rather than quantities of the form $|u^m(p^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m})|$. This is because the difference in unilateral deviations provides a better characterization of the strategic similarities (equilibrium and dynamic properties) between two games than the difference in utility functions, by “quotienting out” the nonstrategic component of games (as discussed in Chapter 2). For this reason, in this chapter we employ MPD to measure closeness of games, as opposed to the norm introduced in Section 2.6.

We conclude this section by presenting a framework for finding the closest potential game to a given game, where the distance between the games is measured in terms of MPD.⁷ Theorem 2.5.1 suggests that the set of potential games is convex, i.e., if $\mathcal{G} = \langle \mathcal{M}, E, \{u^m\}_m \rangle$ and $\hat{\mathcal{G}} = \langle \mathcal{M}, E, \{\hat{u}^m\}_m \rangle$ are potential games, then $\mathcal{G}_\alpha = \langle \mathcal{M}, E, \{\alpha u^m + (1 - \alpha)\hat{u}^m\}_m \rangle$, is also a potential game provided that $\alpha \in [0, 1]$. Hence, intuitively, it should suffice to solve a convex optimization problem to find the closest potential game to a given game. We next provide one such convex optimization formulation.

Assume that a game with utility functions $\{u^m\}_m$ is given. The closest potential game (in terms of MPD) to this game, with payoff functions $\{\hat{u}^m\}_m$, and potential function ϕ can be obtained by solving the following optimization problem:

$$\begin{aligned}
 \min_{\phi, \{\hat{u}^m\}_m} \quad & \max_{\mathbf{p} \in E, m \in \mathcal{M}, q^m \in E^m} | (u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})) \\
 \text{(P :)} \quad & \quad \quad \quad - (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m})) | \\
 \text{s.t.} \quad & \phi(\bar{q}^m, \bar{\mathbf{p}}^{-m}) - \phi(\bar{p}^m, \bar{\mathbf{p}}^{-m}) = \hat{u}^m(\bar{q}^m, \bar{\mathbf{p}}^{-m}) - \hat{u}^m(\bar{p}^m, \bar{\mathbf{p}}^{-m}), \\
 & \text{for all } m \in \mathcal{M}, \bar{\mathbf{p}} \in E, \bar{q}^m \in E^m.
 \end{aligned}$$

Note that the difference $(u^m(q^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m})) - (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m}))$ is linear in $\{\hat{u}^m\}_m$. Thus, the objective function is the maximum of such linear functions, and hence is convex in $\{\hat{u}^m\}_m$. The constraints of this optimization problem guarantee that the game with payoff functions $\{\hat{u}^m\}_m$ is a potential game with potential ϕ . Note that these constraints are linear. Therefore, it follows that (P) is a convex optimization problem that gives the closest potential game to a given game.

In the rest of the chapter, we do not discuss how a close potential game to a given game is obtained, but we just assume that a close potential game with potential ϕ is known and the MPD between this game and the original game is δ . We provide characterization results on limiting dynamics for a given game in terms of ϕ and δ .

3.3 Better Response and Best Response Dynamics

In this section, we consider better and best response dynamics, and study convergence properties of these update rules in near-potential games. All of the update rules considered in

⁷A similar framework for finding near-potential (and weighted potential) games (using a different norm) can be found in Candogan et al. (2011a) and Candogan et al. (2010a).

this section are discrete-time update rules, i.e., players are allowed to update their strategies at time instants $t \in \mathbb{Z}_+ = \{1, 2, \dots\}$.

Best response dynamics is an update rule where at each time instant a player chooses its best response to other players' current strategy profile. In better response dynamics, on the other hand, players choose strategies that improve their payoffs, but these strategies need not be their best responses. Formal descriptions of these update rules are given below.

Definition 3.3.1 (Better and Best Response Dynamics). *At each time instant $t \in \{1, 2, \dots\}$, a single player is chosen at random for updating its strategy, using a probability distribution with full support over the set of players. Let m be the player chosen at some time t , and let $\mathbf{r} \in E$ denote the strategy profile that is used at time $t - 1$.*

1. *Better response dynamics is the update process where player m does not modify its strategy if $u^m(\mathbf{r}) = \max_{q^m} u^m(q^m, \mathbf{r}^{-m})$, and otherwise it updates its strategy to a strategy in $\{q^m | u^m(q^m, \mathbf{r}^{-m}) > u^m(\mathbf{r})\}$, chosen uniformly at random.*
2. *Best response dynamics is the update process where player m does not modify its strategy if $u^m(\mathbf{r}) = \max_{q^m} u^m(q^m, \mathbf{r}^{-m})$, and otherwise it updates its strategy to a strategy in $\arg \max_{q^m} u^m(q^m, \mathbf{r}^{-m})$, chosen uniformly at random.*

For simplicity of the analysis, we assume here that users are chosen randomly to update their strategy. However, this assumption is not crucial for our results, and can be relaxed.

We refer to strategies in $\arg \max_{q^m} u^m(q^m, \mathbf{r}^{-m})$ as *best responses of player m to \mathbf{r}^{-m}* . We denote the strategy profile used at time t by \mathbf{p}_t , and we define the *trajectory of the dynamics* as the sequence of strategy profiles $\{\mathbf{p}_t\}_{t=0}^\infty$. In our analysis, we assume that the trajectory is initialized at a strategy profile $\mathbf{p}_0 \in E$ at time 0 and it evolves according to one of the update rules described above.

The following theorem establishes that in finite games, better and best response dynamics converge to a set of ϵ -equilibria, where the size of this set is characterized by the MPD to a close potential game.⁸

Theorem 3.3.1. *Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a nearby potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Assume that best response or better response dynamics are used in \mathcal{G} , and denote the number of strategy profiles in these games by $|E| = h$.*

For both update processes, the trajectories are contained in the δh -equilibrium set of \mathcal{G} after finite time with probability 1, i.e., let T be a random variable such that $\mathbf{p}_t \in \mathcal{X}_{\delta h}$, for all $t > T$, then $P(T < \infty) = 1$.

Proof. We prove the claim by modeling the update process using a Markov chain, and employing the improvement path condition for potential games (cf. Proposition 3.2.1).

⁸The bound of this theorem can be improved by using bounds on the length of utility improvement cycles in games. In particular, for two player games using such a bound due to Ahn (2006), it is possible to establish convergence to a smaller approximate equilibrium set.

Using Definition 3.3.1, we can represent the strategy updates in best response dynamics as the state transitions in the following Markov chain: (i) Each state corresponds to a strategy profile and, (ii) there is a nonzero transition probability from state \mathbf{r} to state $\mathbf{q} \neq \mathbf{r}$, if \mathbf{r} and \mathbf{q} differ in the strategy of a single player, say m , and q^m is a (strict) best response of player m to \mathbf{r}^{-m} . The probability of transition from state \mathbf{r} to state \mathbf{q} is equal to the probability that at strategy profile \mathbf{r} , player m is chosen for update and it chooses q^m as its new strategy. In the case of better response dynamics we allow q^m to be any strategy strictly improving payoff of player m , and a similar Markov chain representation still holds.

Since there are finitely many states, one of the recurrent classes of the Markov chain is reached in finite time (with probability 1). Thus, to prove the claim, it is sufficient to show that any state which belongs to some recurrent class of this Markov chain is contained in the ϵ -equilibrium set of \mathcal{G} .

It follows from Definition 3.3.1 that a recurrence class is a singleton, only if none of the players can strictly improve its payoff by unilaterally deviating from the corresponding strategy profile. Thus, such a strategy profile is a Nash equilibrium of \mathcal{G} and is contained in the ϵ -equilibrium set.

Consider a recurrence class that is not a singleton. Let \mathbf{r} be a strategy profile in this recurrence class. Since the recurrence class is not a singleton, there exists some player m , who can unilaterally deviate from \mathbf{r} by following its best response to another strategy profile \mathbf{q} , and increase its payoff by some $\alpha > 0$. Since such a transition occurs with nonzero probability, \mathbf{r} and \mathbf{q} are in the same recurrence class, and the process when started from \mathbf{r} visits \mathbf{q} and returns to \mathbf{r} in finitely many updates. Since each transition corresponds to a unilateral deviation that strictly improves the payoff of the deviating player, this constitutes a simple closed improvement path containing \mathbf{r} and \mathbf{q} . Let $\gamma = (\mathbf{p}_0, \dots, \mathbf{p}_N)$ be such an improvement path and $\mathbf{p}_0 = \mathbf{p}_N = \mathbf{r}$, $\mathbf{p}_1 = \mathbf{q}$ and $N \leq |E| = h$. Since $u^m(\mathbf{q}) - u^m(\mathbf{r}) = \alpha$, and $u^{m_i}(\mathbf{p}_i) - u^{m_i}(\mathbf{p}_{i-1}) \geq 0$ at every step i of the path, this closed improvement path satisfies

$$\sum_{i=1}^N (u^{m_i}(\mathbf{p}_i) - u^{m_i}(\mathbf{p}_{i-1})) \geq \alpha. \quad (3.1)$$

On the other hand it follows by Proposition 3.2.1 that the close potential game satisfies

$$\sum_{i=1}^N (\hat{u}^{m_i}(\mathbf{p}_i) - \hat{u}^{m_i}(\mathbf{p}_{i-1})) = 0. \quad (3.2)$$

Combining (3.1) and (3.2) we conclude that

$$\begin{aligned} \alpha &\leq \sum_{i=1}^N (u^{m_i}(\mathbf{p}_i) - u^{m_i}(\mathbf{p}_{i-1})) - (\hat{u}^{m_i}(\mathbf{p}_i) - \hat{u}^{m_i}(\mathbf{p}_{i-1})) \\ &\leq N\delta. \end{aligned}$$

Since $N \leq |E| = h$, it follows that $\alpha \leq \delta h$. The claim then immediately follows since \mathbf{r} and the recurrence class were chosen arbitrarily, and our analysis shows that the payoff improvement of player m (chosen for strategy update using a probability distribution with full support as described in Definition 3.3.1), due to its best response is bounded by δh . \square

As can be seen from the proof of this theorem, extending dynamical properties of potential games to nearby games relies on special structural properties of potential games. As a corollary of the above theorem, we obtain that trajectories generated by better and best response dynamics converge to a Nash equilibrium in potential games, since if \mathcal{G} is a potential game, the close potential game $\hat{\mathcal{G}}$ can be chosen such that $d(\mathcal{G}, \hat{\mathcal{G}}) = 0$. Also, it follows from our proof that our result is applicable in cases where the underlying game has better response cycles. Thus, even when the game does not share similar ordinal properties with potential games, our approach can be used to approximately characterize the limiting behavior of dynamics.

3.4 Logit Response Dynamics

In this section we focus on logit response dynamics. Logit response dynamics can be viewed as a smoothed version of the best response dynamics, in which a smoothing parameter determines the frequency with which the best response strategy is picked. The evolution of the pure strategy profiles can be represented in terms of a Markov chain (with state space given by the set of pure strategy profiles). We characterize the stationary distribution and stochastically stable states of this Markov chain (or of the update rule) in near-potential games. Our approach involves identifying a close potential game to a given game, and exploiting features of the corresponding potential function to characterize the limiting behavior of logit response dynamics in the original game.

In Section 3.4.1, we provide a formal definition of logit response dynamics and review some of its properties. We also present some of the mathematical tools used in the literature to study this update rule. In Section 3.4.2, we show that the stationary distribution of logit response dynamics in a near-potential game can be approximately characterized using the potential function of a nearby potential game. We also use this result to show that the stochastically stable strategy profiles are contained in approximate equilibrium sets in near-potential games.

3.4.1 Properties of Logit Response

We start by providing a formal definition of logit response dynamics:

Definition 3.4.1. *At each time instant $t \in \{1, 2, \dots\}$, a single player is chosen at random for updating its strategy, using a probability distribution with full support over the set of*

players. Let m be the player chosen at some time t , and let $\mathbf{r} \in E$ denote the strategy profile that is used at time $t - 1$.

Logit response dynamics with parameter τ is the update process, where player m chooses a strategy $q^m \in E^m$ with probability

$$P_\tau^m(q^m|\mathbf{r}) = \frac{e^{\frac{1}{\tau}u^m(q^m, \mathbf{r}^{-m})}}{\sum_{p^m \in E^m} e^{\frac{1}{\tau}u^m(p^m, \mathbf{r}^{-m})}}.$$

In this definition, $\tau > 0$ is a fixed parameter that determines how often players choose their best responses. The probability of not choosing a best response decreases as τ decreases, and as $\tau \rightarrow 0$, players choose their best responses with probability 1. This feature suggests that logit response dynamics can be viewed as a generalization of best response dynamics, where with small but nonzero probability players use a strategy that is not a best response.

For a given $\tau > 0$, this update process can be represented by a finite aperiodic and irreducible Markov chain (Marden and Shamma, 2012; Alós-Ferrer and Netzer, 2010). The states of the Markov chain correspond to the strategy profiles in the game. Denoting the probability that player m is chosen for a strategy update by α_m , transition probability from strategy profile \mathbf{p} to \mathbf{q} can be given by (assuming $\mathbf{p} \neq \mathbf{q}$, and denoting the transition from \mathbf{p} to \mathbf{q} by $\mathbf{p} \rightarrow \mathbf{q}$):

$$P_\tau(\mathbf{p} \rightarrow \mathbf{q}) = \begin{cases} \alpha_m P_\tau^m(q^m|\mathbf{p}) & \text{if } \mathbf{q}^{-m} = \mathbf{p}^{-m} \text{ for some } m \in \mathcal{M} \\ 0 & \text{otherwise.} \end{cases} \quad (3.3)$$

The chain is aperiodic and irreducible since a player updating its strategy can choose any strategy (including the current one) with positive probability. Consequently, it has a unique stationary distribution.

We denote the stationary distribution of this Markov chain by μ_τ and refer to it as the stationary distribution of the logit response dynamics. A strategy profile \mathbf{q} such that $\lim_{\tau \rightarrow 0} \mu_\tau(\mathbf{q}) > 0$ is referred to as a *stochastically stable strategy profile* of the logit response dynamics. Intuitively, these strategy profiles are the ones that are used with nonzero probability, as players adopt their best responses more and more frequently in their strategy updates.

In potential games, the stationary distribution of the logit response dynamics can be written as an explicit function of the potential. If \mathcal{G} is a potential game with potential function ϕ , the stationary distribution of the logit response dynamics is given by the distribution

(Marden and Shamma, 2012; Alós-Ferrer and Netzer, 2010; Blume, 1997):⁹

$$\mu_\tau(\mathbf{q}) = \frac{e^{\frac{1}{\tau}\phi(\mathbf{q})}}{\sum_{\mathbf{p} \in E} e^{\frac{1}{\tau}\phi(\mathbf{p})}}. \quad (3.4)$$

It can be seen from (3.4) that $\lim_{\tau \rightarrow 0} \mu_\tau(\mathbf{q}) > 0$ if and only if $\mathbf{q} \in \arg \max_{\mathbf{p} \in E} \phi(\mathbf{p})$. Thus, in potential games the stochastically stable strategy profiles are those that maximize the potential function.

We next describe a method for obtaining the stationary distribution of Markov chains. This method will be used in the next subsection in characterizing the stationary distribution of logit response. Assume that an irreducible Markov chain over a finite set of states S , with transition probability matrix P is given. Consider a directed tree, T , with nodes given by the states of the Markov chain, and assume that an edge from node \mathbf{q} to node \mathbf{p} can exist only if there is a nonzero transition probability from \mathbf{q} to \mathbf{p} in the Markov chain. We say that the tree is rooted at state \mathbf{p} , if from every state $\mathbf{q} \neq \mathbf{p}$ there exists a unique directed path along the tree to \mathbf{p} . For each state $\mathbf{p} \in S$, denote by $\mathcal{T}(\mathbf{p})$ the set of all trees rooted at \mathbf{p} , and define a weight $w_{\mathbf{p}} \geq 0$ such that

$$w_{\mathbf{p}} = \sum_{T \in \mathcal{T}(\mathbf{p})} \prod_{(\mathbf{q} \rightarrow \mathbf{r}) \in T} P(\mathbf{q} \rightarrow \mathbf{r}). \quad (3.5)$$

The following proposition from the Markov Chain literature (Leighton and Rivest (1983); Anantharam and Tsoucas (1989); Freidlin and Wentzell (1998)), known as the Markov chain tree theorem, expresses the stationary distribution of Markov chains in terms of these weights.

Proposition 3.4.1. *The stationary distribution of the Markov chain defined over set S is given by $\mu(\mathbf{p}) = \frac{w_{\mathbf{p}}}{\sum_{\mathbf{q} \in S} w_{\mathbf{q}}}$.*

For any $T \in \mathcal{T}(\mathbf{p})$, intuitively, the quantity $\prod_{(\mathbf{q} \rightarrow \mathbf{r}) \in T} P(\mathbf{q} \rightarrow \mathbf{r})$ gives a measure of likelihood of the event that node \mathbf{p} is reached when the chain is initiated from the leaves (i.e., nodes with indegree equal to 0) of T . Thus, $w_{\mathbf{p}}$ captures how likely it is that node \mathbf{p} is visited in this chain, and the normalization in Proposition 3.4.1 gives the stationary distribution. Since for finite games logit response dynamics can be modeled as an irreducible Markov chain, this result can be used to characterize its stationary distribution.

3.4.2 Stationary Distribution of Logit Response Dynamics

In this section we show that the stationary distribution of logit response dynamics in near-potential games can be approximated by exploiting the potential function of a close potential

⁹Note that this expression is independent of $\{\alpha_m\}$, i.e., the probability distribution that is used to choose which player updates its strategy has no effect on the stationary distribution of logit response.

game. We start by showing that in games with small MPD logit response dynamics have similar transition probabilities.

Lemma 3.4.1. *Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a nearby potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the transition probability matrices of logit response dynamics in \mathcal{G} and $\hat{\mathcal{G}}$ by P_τ and \hat{P}_τ respectively. For all strategy profiles \mathbf{p} and \mathbf{q} that differ in the strategy of at most one player, we have*

$$e^{-\frac{2\delta}{\tau}} \leq \hat{P}_\tau(\mathbf{p} \rightarrow \mathbf{q})/P_\tau(\mathbf{p} \rightarrow \mathbf{q}) \leq e^{\frac{2\delta}{\tau}}.$$

Proof. Assume that $\mathbf{p}^{-m} = \mathbf{q}^{-m}$. In \mathcal{G} the transition probability $P_\tau(\mathbf{p} \rightarrow \mathbf{q})$ can be expressed by (see (3.3)):

$$P_\tau(\mathbf{p} \rightarrow \mathbf{q}) = \begin{cases} \alpha_m P_\tau^m(q^m | \mathbf{p}) & \text{if } q^m \neq p^m \\ \sum_{k \in \mathcal{M}} \alpha_k P_\tau^k(p^k | \mathbf{p}) & \text{otherwise.} \end{cases}$$

A similar expression holds for the transition probability $\hat{P}_\tau(\mathbf{p} \rightarrow \mathbf{q})$ in $\hat{\mathcal{G}}$, replacing P_τ^m by \hat{P}_τ^m . Thus, it is sufficient to prove $e^{-\frac{2\delta}{\tau}} \leq \hat{P}_\tau^m(q^m | \mathbf{p})/P_\tau^m(q^m | \mathbf{p}) \leq e^{\frac{2\delta}{\tau}}$ for all \mathbf{p} , m , q^m to prove the claim.

Observe that by the definition of MPD

$$\begin{aligned} u^m(r^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m}) - \delta &\leq \hat{u}^m(r^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m}) \\ &\leq u^m(r^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m}) + \delta. \end{aligned} \quad (3.6)$$

Definition 3.4.1 suggests that $\hat{P}_\tau^m(q^m | \mathbf{p})$ can be written as (by dividing the numerator and the denominator by $e^{\frac{1}{\tau} \hat{u}^m(p^m, \mathbf{p}^{-m})}$):

$$\hat{P}_\tau^m(q^m | \mathbf{p}) = \frac{e^{\frac{1}{\tau} (\hat{u}^m(q^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m}))}}{\sum_{r^m \in E^m} e^{\frac{1}{\tau} (\hat{u}^m(r^m, \mathbf{p}^{-m}) - \hat{u}^m(p^m, \mathbf{p}^{-m}))}}.$$

Therefore, using the bounds in (3.6) it follows that

$$\hat{P}_\tau^m(q^m | \mathbf{p}) \leq \frac{\kappa(q^m) e^{\frac{\delta}{\tau}}}{\kappa(q^m) e^{\frac{\delta}{\tau}} + \sum_{r^m \neq q^m} \kappa(r^m) e^{-\frac{\delta}{\tau}}}.$$

where, $\kappa(r^m) = e^{\frac{1}{\tau} (u^m(r^m, \mathbf{p}^{-m}) - u^m(p^m, \mathbf{p}^{-m}))}$ for all $r^m \in E^m$. Dividing both the numerator and the denominator of the right hand side by $\sum_{r^m \in E^m} \kappa(r^m)$ and observing that $P_\tau^m(q^m | \mathbf{p}) = \frac{\kappa(q^m)}{\sum_{r^m \in E^m} \kappa(r^m)}$, we obtain

$$\hat{P}_\tau^m(q^m | \mathbf{p}) \leq \frac{e^{\frac{\delta}{\tau}} P_\tau^m(q^m | \mathbf{p})}{e^{\frac{\delta}{\tau}} P_\tau^m(q^m | \mathbf{p}) + e^{-\frac{\delta}{\tau}} (1 - P_\tau^m(q^m | \mathbf{p}))},$$

or equivalently

$$\frac{\hat{P}_\tau^m(q^m|\mathbf{p})}{P_\tau^m(q^m|\mathbf{p})} \leq \frac{e^{\frac{\delta}{\tau}}}{e^{\frac{\delta}{\tau}}P_\tau^m(q^m|\mathbf{p}) + e^{-\frac{\delta}{\tau}}(1 - P_\tau^m(q^m|\mathbf{p}))}.$$

It can be seen that the right hand side is decreasing in $P_\tau^m(q^m|\mathbf{p})$. Thus replacing $P_\tau^m(q^m|\mathbf{p})$ by 0, the right hand side can be upper bounded by $e^{\frac{2\delta}{\tau}}$. Then we obtain $\hat{P}_\tau^m(q^m|\mathbf{p})/P_\tau^m(q^m|\mathbf{p}) \leq e^{\frac{2\delta}{\tau}}$. By symmetry we also conclude that $P_\tau^m(q^m|\mathbf{p})/\hat{P}_\tau^m(q^m|\mathbf{p}) \leq e^{\frac{2\delta}{\tau}}$, and combining these bounds the claim follows. \square

Definition 3.4.1 suggests that perturbation of utility functions changes the transition probabilities multiplicatively in logit response. The above lemma supports this intuition: if utility gains due to unilateral deviations are modified by δ , the ratio of the transition probabilities can change at most by $e^{\frac{2\delta}{\tau}}$. Thus, if two games are close, then the transition probabilities of logit response in these games should be closely related.

This suggests using results from perturbation theory of Markov chains to characterize the stationary distribution of logit response in a near-potential game (Haviv and Van der Heyden, 1984; Cho and Meyer, 2001). However, standard perturbation results characterize changes in the stationary distribution of a Markov chain when the transition probabilities are *additively perturbed*. These results, when applied to multiplicative perturbations, yield bounds which are uninformative. We therefore first present a result which characterizes deviations from the stationary distribution of a Markov chain when its transition probabilities are multiplicatively perturbed, and therefore may be of independent interest.¹⁰

Theorem 3.4.1. *Let P and \hat{P} denote the probability transition matrices of two finite irreducible Markov chains with the same state space. Denote the stationary distributions of these Markov chains by μ and $\hat{\mu}$ respectively, and let the cardinality of the state space be h . Assume that $\alpha \geq 1$ is a given constant and for any two states \mathbf{p} and \mathbf{q} , the following inequalities hold*

$$\alpha^{-1}P(\mathbf{p} \rightarrow \mathbf{q}) \leq \hat{P}(\mathbf{p} \rightarrow \mathbf{q}) \leq \alpha P(\mathbf{p} \rightarrow \mathbf{q}).$$

Then, for any state \mathbf{p} , we have

$$(i) \quad \frac{\alpha^{-(h-1)}\mu(\mathbf{p})}{\alpha^{-(h-1)}\mu(\mathbf{p}) + \alpha^{h-1}(1 - \mu(\mathbf{p}))} \leq \hat{\mu}(\mathbf{p}) \leq \frac{\alpha^{h-1}\mu(\mathbf{p})}{\alpha^{h-1}\mu(\mathbf{p}) + \alpha^{-(h-1)}(1 - \mu(\mathbf{p}))},$$

$$(ii) \quad |\mu(\mathbf{p}) - \hat{\mu}(\mathbf{p})| \leq \frac{\alpha^{h-1} - 1}{\alpha^{h-1} + 1}.$$

Proof. As before, let $\mathcal{T}(\mathbf{p})$ denote the set of directed trees that are rooted at state \mathbf{p} . Using the characterization of the stationary distribution in Proposition 3.4.1, for the Markov chain

¹⁰A multiplicative perturbation bound similar to ours, can be found in Freidlin and Wentzell (1998). However, this bound is looser than the one we obtain and it does not provide a good characterization of the stationary distribution in our setting. We provide a tighter bound, and obtain stronger predictions on the stationary distribution of logit response.

with probability transition matrix P , we have $\mu(\mathbf{p}) = \frac{w_{\mathbf{p}}}{\sum_{\mathbf{q}} w_{\mathbf{q}}}$, where for each state \mathbf{p} ,

$$w_{\mathbf{p}} = \sum_{T \in \mathcal{T}(\mathbf{p})} \prod_{(\mathbf{x} \rightarrow \mathbf{y}) \in T} P(\mathbf{x} \rightarrow \mathbf{y}).$$

For the Markov chain with probability transition matrix \hat{P} , we define $\hat{w}_{\mathbf{p}}$, by replacing P in the above equation with \hat{P} and $\hat{\mu}(\mathbf{p})$ similarly satisfies $\hat{\mu}(\mathbf{p}) = \frac{\hat{w}_{\mathbf{p}}}{\sum_{\mathbf{q}} \hat{w}_{\mathbf{q}}}$.

Since the Markov chain has h states, $|T| = h - 1$ for all $T \in \mathcal{T}(\mathbf{p})$. Hence, it follows from the assumption of the theorem and the above definitions of $w_{\mathbf{p}}$ and $\hat{w}_{\mathbf{p}}$ that

$$\begin{aligned} \alpha^{-(h-1)} w_{\mathbf{p}} &= \alpha^{-(h-1)} \sum_{T \in \mathcal{T}(\mathbf{p})} \prod_{(\mathbf{x} \rightarrow \mathbf{y}) \in T} P(\mathbf{x} \rightarrow \mathbf{y}) \\ &\leq \hat{w}_{\mathbf{p}} = \sum_{T \in \mathcal{T}(\mathbf{p})} \prod_{(\mathbf{x} \rightarrow \mathbf{y}) \in T} \hat{P}(\mathbf{x} \rightarrow \mathbf{y}) \\ &\leq \alpha^{h-1} \sum_{T \in \mathcal{T}(\mathbf{p})} \prod_{(\mathbf{x} \rightarrow \mathbf{y}) \in T} P(\mathbf{x} \rightarrow \mathbf{y}) = \alpha^{h-1} w_{\mathbf{p}}. \end{aligned}$$

This inequality implies that for all \mathbf{q} , $\hat{w}_{\mathbf{q}}$ is upper bounded by $\alpha^{h-1} w_{\mathbf{q}}$ and lower bounded by $\alpha^{-(h-1)} w_{\mathbf{q}}$. Using this observation together with the identity $\hat{\mu}(\mathbf{p}) = \frac{\hat{w}_{\mathbf{p}}}{\sum_{\mathbf{q}} \hat{w}_{\mathbf{q}}}$, we obtain

$$\frac{\alpha^{-(h-1)} w_{\mathbf{p}}}{\alpha^{-(h-1)} w_{\mathbf{p}} + \alpha^{h-1} \sum_{\mathbf{q} \neq \mathbf{p}} w_{\mathbf{q}}} \leq \hat{\mu}(\mathbf{p}) = \frac{\hat{w}_{\mathbf{p}}}{\sum_{\mathbf{q}} \hat{w}_{\mathbf{q}}} \leq \frac{\alpha^{h-1} w_{\mathbf{p}}}{\alpha^{h-1} w_{\mathbf{p}} + \alpha^{-(h-1)} \sum_{\mathbf{q} \neq \mathbf{p}} w_{\mathbf{q}}}.$$

Dividing the numerators and denominators of the left and right hand sides of the inequality by $\sum_{\mathbf{q}} w_{\mathbf{q}}$, using Proposition 3.4.1, and observing that $\sum_{\mathbf{q} \neq \mathbf{p}} \mu(\mathbf{q}) = 1 - \mu(\mathbf{p})$ the first part of the theorem follows.

Consider functions f and g defined on $[0, 1]$ such that $f(x) = \frac{\alpha^{h-1} x}{\alpha^{h-1} x + \alpha^{-(h-1)}(1-x)} - x$ and $g(x) = \frac{\alpha^{-(h-1)} x}{\alpha^{-(h-1)} x + \alpha^{h-1}(1-x)} - x$ for $x \in [0, 1]$. Checking the first order optimality conditions, it can be seen that $f(x)$ is maximized at $x = \frac{\alpha^{-(h-1)}}{1 + \alpha^{-(h-1)}}$, and the maximum equals to $\frac{\alpha^{h-1} - 1}{\alpha^{h-1} + 1}$. Similarly, the minimum of $g(x)$ is achieved at $x = \frac{\alpha^{h-1}}{1 + \alpha^{h-1}}$ and is equal to $\frac{1 - \alpha^{h-1}}{1 + \alpha^{h-1}}$. Combining these observations with part (i), we obtain

$$\begin{aligned} \frac{1 - \alpha^{h-1}}{1 + \alpha^{h-1}} \leq g(\mu(\mathbf{p})) &= \frac{\alpha^{-(h-1)} \mu(\mathbf{p})}{\alpha^{-(h-1)} \mu(\mathbf{p}) + \alpha^{h-1} (1 - \mu(\mathbf{p}))} - \mu(\mathbf{p}) \leq \hat{\mu}(\mathbf{p}) - \mu(\mathbf{p}) \\ &\leq \frac{\alpha^{h-1} \mu(\mathbf{p})}{\alpha^{h-1} \mu(\mathbf{p}) + \alpha^{-(h-1)} (1 - \mu(\mathbf{p}))} - \mu(\mathbf{p}) = f(\mu(\mathbf{p})) \leq \frac{\alpha^{h-1} - 1}{\alpha^{h-1} + 1}, \end{aligned}$$

hence the second part of the claim follows. \square

Next we use the above theorem to relate the stationary distributions of logit response dynamics in nearby games.

Corollary 3.4.1. *Let \mathcal{G} and $\hat{\mathcal{G}}$ be finite games with number of strategy profiles $|E| = h$, such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the stationary distributions of logit response dynamics in these games by μ_τ , and $\hat{\mu}_\tau$ respectively. Then, for any strategy profile \mathbf{p} we have*

$$(i) \quad \frac{e^{-\frac{2\delta(h-1)}{\tau}} \mu_\tau(\mathbf{p})}{e^{-\frac{2\delta(h-1)}{\tau}} \mu_\tau(\mathbf{p}) + e^{\frac{2\delta(h-1)}{\tau}} (1 - \mu_\tau(\mathbf{p}))} \leq \hat{\mu}_\tau(\mathbf{p}) \leq \frac{e^{\frac{2\delta(h-1)}{\tau}} \mu_\tau(\mathbf{p})}{e^{\frac{2\delta(h-1)}{\tau}} \mu_\tau(\mathbf{p}) + e^{-\frac{2\delta(h-1)}{\tau}} (1 - \mu_\tau(\mathbf{p}))},$$

$$(ii) \quad |\mu_\tau(\mathbf{p}) - \hat{\mu}_\tau(\mathbf{p})| \leq \frac{e^{\frac{2\delta(h-1)}{\tau}} - 1}{e^{\frac{2\delta(h-1)}{\tau}} + 1}.$$

Proof. Proof follows from Lemma 3.4.1 and Theorem 3.4.1 by setting $\alpha = e^{\frac{2\delta}{\tau}}$. \square

The above corollary can be adapted to near-potential games, by exploiting the relation of stationary distribution of logit response and potential function in potential games (see (3.4)). We conclude this section by providing such a characterization of the stationary distribution of logit response dynamics in near-potential games.

Corollary 3.4.2. *Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a nearby potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the potential function of $\hat{\mathcal{G}}$ by ϕ , and the number of strategy profiles in these games by $|E| = h$. Then, the stationary distribution μ_τ of logit response dynamics in \mathcal{G} is such that*

$$(i) \quad \frac{e^{\frac{1}{\tau}(\phi(\mathbf{p}) - 2\delta(h-1))}}{e^{\frac{1}{\tau}(\phi(\mathbf{p}) - 2\delta(h-1))} + \sum_{\mathbf{q} \neq \mathbf{p} \in E} e^{\frac{1}{\tau}(\phi(\mathbf{q}) + 2\delta(h-1))}} \leq \mu_\tau(\mathbf{p})$$

$$\leq \frac{e^{\frac{1}{\tau}(\phi(\mathbf{p}) + 2\delta(h-1))}}{e^{\frac{1}{\tau}(\phi(\mathbf{p}) + 2\delta(h-1))} + \sum_{\mathbf{q} \neq \mathbf{p} \in E} e^{\frac{1}{\tau}(\phi(\mathbf{q}) - 2\delta(h-1))}},$$

$$(ii) \quad \left| \mu_\tau(\mathbf{p}) - \frac{e^{\frac{1}{\tau}\phi(\mathbf{p})}}{\sum_{\mathbf{q} \in E} e^{\frac{1}{\tau}\phi(\mathbf{q})}} \right| \leq \frac{e^{\frac{2\delta(h-1)}{\tau}} - 1}{e^{\frac{2\delta(h-1)}{\tau}} + 1}.$$

Proof. Proof follows from Corollary 3.4.1 and (3.4). \square

With simple manipulations, it can be shown that $(e^x - 1)/(e^x + 1) \leq x/2$ for $x \geq 0$. Thus, (ii) in the above corollary implies that $\left| \mu_\tau(\mathbf{p}) - \frac{e^{\frac{1}{\tau}\phi(\mathbf{p})}}{\sum_{\mathbf{q} \in E} e^{\frac{1}{\tau}\phi(\mathbf{q})}} \right| \leq \frac{\delta(h-1)}{\tau}$. Therefore, the stationary distribution of logit response dynamics in a near-potential game can be characterized in terms of the stationary distribution of this update rule in a close potential game. When τ is fixed and $\delta \rightarrow 0$, i.e., when the original game is arbitrarily close to a potential game, the stationary distribution of logit response is arbitrarily close to the stationary distribution in the potential game. On the other hand, for a fixed δ , as $\tau \rightarrow 0$, the upper bound in (ii) becomes uninformative. This is the case since $\tau \rightarrow 0$ implies that players adopt their best responses with probability 1, and thus the stationary distribution of the update rule becomes very sensitive to the difference of the game from a potential

game. In this case we can still characterize the stochastically stable states of logit response using the results of Corollary 3.4.2, as we show in Corollary 3.4.3.

Corollary 3.4.3. *Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a nearby potential game with potential function ϕ and $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the potential function of $\hat{\mathcal{G}}$ by $\hat{\phi}$, and the number of strategy profiles in these games by $|E| = h$. The stochastically stable strategy profiles of \mathcal{G} are (i) contained in $S = \{\mathbf{p} | \phi(\mathbf{p}) \geq \max_{\mathbf{q}} \phi(\mathbf{q}) - 4\delta(h-1)\}$, (ii) $4\delta h$ -equilibria of \mathcal{G} .*

Proof. (i) The upper bound in the first part of Corollary 3.4.2 implies that if \mathbf{p} is a strategy profile such that $\phi(\mathbf{p}) < \max_{\mathbf{q} \in E} \phi(\mathbf{q}) - 4\delta(h-1)$, then the stationary distribution of logit response in \mathcal{G} is such that $\mu_\tau(\mathbf{p}) \rightarrow 0$ as $\tau \rightarrow 0$. Thus, it immediately follows that the stochastically stable states in \mathcal{G} are contained in $\{\mathbf{p} \in E | \phi(\mathbf{p}) \geq \max_{\mathbf{q} \in E} \phi(\mathbf{q}) - 4\delta(h-1)\}$.

(ii) From the definition of S it follows that in $\hat{\mathcal{G}}$, none of the players can deviate from a strategy profile in S and improve its utility by more than $4\delta(h-1)$. Since $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$ it follows from part (i) that in \mathcal{G} , none of the players can unilaterally deviate from a stochastically stable strategy profile and improve its utility by more than $4\delta(h-1) + \delta \leq 4\delta h$. Hence stochastically stable strategy profiles of \mathcal{G} are $4\delta h$ -equilibria. \square

We conclude that in near-potential games, the stochastically stable states of logit response are the strategy profiles that approximately maximize the potential function of a close potential game. This result enables us to characterize the stochastically stable states of logit response dynamics in near-potential games, without explicitly computing the stationary distribution.

Since it is possible to identify a potential game that is close to a given game (as explained in Section 3.2), Corollaries 3.4.2 and 3.4.3 provide a systematic approach for characterizing the stationary distribution and stochastically stable states of logit response, for general games. The characterization is tighter for near-potential games, but it is still informative for general games.

Moreover, our results enable robust predictions about stochastically stable strategy profiles in potential games. In particular, we can quantify payoff perturbations that maintain stochastically stable states of a game. For instance, consider potential games where the potential ϕ has a unique maximizer \mathbf{q}^* . Corollary 3.4.3 implies that in such games if the payoffs are perturbed by at most $\frac{1}{8h} (\phi(\mathbf{q}^*) - \max_{\mathbf{q} \neq \mathbf{q}^*} \phi(\mathbf{q}))$ (so that the MPD between the original game and the game obtained after perturbations satisfies $\delta \leq \frac{1}{4h} (\phi(\mathbf{q}^*) - \max_{\mathbf{q} \neq \mathbf{q}^*} \phi(\mathbf{q}))$) the stochastically stable strategy profiles do not change. Thus, the corollary allows us to bound the magnitude of payoff perturbations that leave the set of stochastically stable strategy profiles intact.

3.5 Fictitious Play

In this section, we investigate the convergence behavior of fictitious play in near-potential games. Unlike better/best response dynamics and logit response, in fictitious play agents maintain an empirical frequency distribution of other players' strategies and play a best response against it. Thus, analyzing fictitious play dynamics requires the notion of mixed strategies and some additional definitions that are provided in Section 3.5.1. In Section 3.5.2 we show that in finite games the empirical frequencies of fictitious play converge to a set which can be characterized in terms of the approximate equilibrium set of the game and the level sets of the potential function of a close potential game. When the original game is sufficiently close to a potential game, we strengthen this result and establish that the empirical frequencies converge to a small neighborhood of mixed equilibria of the game, and the size of this neighborhood is a function of the distance of the original game from a potential game. As a special case, our result allows us to recover the result of Monderer and Shapley (1996a), which states that in potential games the empirical frequencies of fictitious play converge to the set of mixed Nash equilibria.

3.5.1 Mixed Strategies and Equilibria

In this section, we introduce some additional notation and definitions, which will be used in Section 3.5.2 when studying convergence properties of fictitious play in near-potential games.

As in Section 2.5.2, we use ΔE^m the set of mixed strategies of player m , and $\prod_{m \in \mathcal{M}} \Delta E^m$ as the set of mixed strategy profiles in the game. For $\mathbf{x} \in \prod_{m \in \mathcal{M}} \Delta E^m$ we denote by $\{u^m(\mathbf{x})\}$ the mixed extension of the utility of player m (as defined in Section 2.5.2), i.e.,

$$u^m(\mathbf{x}) = \sum_{\mathbf{p} \in E} u^m(\mathbf{p}) \prod_{k \in \mathcal{M}} x^k(p^k). \quad (3.7)$$

Similarly, we denote the mixed extension of the potential function by $\phi(\mathbf{x})$. We use $\|\cdot\|$ to denote the standard 2-norm on $\prod_m \Delta E^m$, i.e., for $\mathbf{x} \in \prod_m \Delta E^m$, we have $\|\mathbf{x}\|^2 = \sum_{m \in \mathcal{M}} \sum_{p^m \in E^m} (x^m(p^m))^2$.

A mixed strategy profile $\mathbf{x} = \{x^m\}_{m \in \mathcal{M}} \in \prod_m \Delta E^m$ is a *mixed ϵ -equilibrium* if for all $m \in \mathcal{M}$ and $p^m \in E^m$,

$$u^m(p^m, \mathbf{x}^{-m}) - u^m(x^m, \mathbf{x}^{-m}) \leq \epsilon. \quad (3.8)$$

Note that if the inequality holds for $\epsilon = 0$, then \mathbf{x} is a mixed Nash equilibrium of the game. In the rest of the chapter, we use the notation \mathcal{X}_ϵ to denote the set of mixed ϵ -equilibria.

Our characterization of the limiting mixed strategy set of fictitious play depends on the number of players in the game. We use $M = |\mathcal{M}|$ as a short-hand notation for this number.

We conclude this section with two technical lemmas which summarize some properties of mixed equilibria and mixed extensions of potential and utility functions. Proofs of these

lemmas can be found in (Candogan et al., 2011b).

The first lemma establishes the Lipschitz continuity of the mixed extensions of the payoff functions and the potential function. It also shows a natural implication of continuity: for any $\epsilon' > \epsilon$, a small enough neighborhood of the ϵ -equilibrium set is contained in the ϵ' -equilibrium set.

Lemma 3.5.1. (i) *Let $\nu : \prod_{m \in \mathcal{M}} E^m \rightarrow \mathbb{R}$ be a mapping from pure strategy profiles to real numbers. Its mixed extension is Lipschitz continuous with a Lipschitz constant of $M \sum_{p \in E} |\nu(\mathbf{p})|$ over the domain $\prod_{m \in \mathcal{M}} \Delta E^m$.*

(ii) *Let $\alpha \geq 0$ and $\gamma > 0$ be given. There exists a small enough $\theta > 0$ such that for any $\|\mathbf{x} - \mathbf{y}\| < \theta$ if $\mathbf{x} \in \mathcal{X}_\alpha$, then $\mathbf{y} \in \mathcal{X}_{\alpha+\gamma}$.*

Lipschitz continuity follows from the fact that mixed extensions are multilinear functions (3.7), with bounded domains. The proof of the second part immediately follows from the Lipschitz continuity of mixed extensions of payoff functions and the definition of approximate equilibria (3.8). Note that the second part implies that for any $\epsilon' > 0$, there exists a small enough neighborhood of equilibria that is contained in the ϵ' -equilibrium set of the game.

We next study the continuity properties of the approximate equilibrium mapping. We first provide the relevant definitions (see Berge (1963); Fudenberg and Tirole (1991)).

Definition 3.5.1 (Upper Semicontinuous Function). *A function $g : X \rightarrow Y \subset \mathbb{R}$ is upper semicontinuous at x_* , if, for each $\epsilon > 0$ there exists a neighborhood U of x_* , such that $g(x) < g(x_*) + \epsilon$ for all $x \in U$. We say g is upper semicontinuous, if it is upper semicontinuous at every point in its domain.*

Alternatively, g is upper semicontinuous if $\limsup_{x_n \rightarrow x_} g(x_n) \leq g(x_*)$ for every x_* in its domain.*

Definition 3.5.2 (Upper Semicontinuous Correspondence). *A correspondence $g : X \rightrightarrows Y$ is upper semicontinuous at x_* , if for any open neighborhood V of $g(x_*)$ there exists a neighborhood U of x_* such that $g(x) \subset V$ for all $x \in U$. We say g is upper semicontinuous, if it is upper semicontinuous at every point in its domain and $g(x)$ is a compact set for each $x \in X$.*

Alternatively, when Y is compact, g is upper semicontinuous if its graph is closed, i.e., the set $\{(x, y) | x \in X, y \in g(x)\}$ is closed.

We next establish upper semicontinuity of the approximate equilibrium mapping.¹¹

¹¹ Here we fix the game, and discuss upper semicontinuity with respect to the ϵ parameter characterizing the ϵ -equilibrium set. We note that this is different than the common results in the literature which discuss upper semicontinuity of the equilibrium set with respect to changes in the utility functions of the underlying game (see Fudenberg and Tirole (1991)).

Lemma 3.5.2. (i) Let $\nu : \prod_{m \in \mathcal{M}} \Delta E^m \rightarrow \mathbb{R}$ be an upper semicontinuous function. The correspondence $g : \mathbb{R} \rightrightarrows \prod_{m \in \mathcal{M}} \Delta E^m$ such that $g(v) = \{\mathbf{x} | \nu(\mathbf{x}) \geq -v\}$ is upper semicontinuous.

(ii) Let $g : \mathbb{R} \rightrightarrows \prod_{m \in \mathcal{M}} \Delta E^m$ be the correspondence such that $g(\alpha) = \mathcal{X}_\alpha$. This correspondence is upper semicontinuous.

Upper semicontinuity of the approximate equilibrium mapping implies that for any given neighborhood of the ϵ -equilibrium set, there exists an $\epsilon' > \epsilon$ such that ϵ' -equilibrium set is contained in this neighborhood. In particular, this implies that every neighborhood of equilibria of the game contains an ϵ' -equilibrium set for some $\epsilon' > 0$. Hence, if disjoint neighborhoods of equilibria are chosen (assuming there are finitely many equilibria), this implies that there exists some $\epsilon' > 0$, such that the ϵ' -equilibrium set is contained in disjoint neighborhoods of equilibria. In the next section, we use this observation to establish convergence of fictitious play to small neighborhoods of equilibria of near-potential games.

3.5.2 Discrete-Time Fictitious Play

Fictitious play is a classical update rule studied in the learning in games literature. In this section, we consider the fictitious play dynamics, proposed in Brown (1951), and explain how the limiting behavior of this dynamical process can be characterized in near-potential games. In particular, we show that the empirical frequencies of fictitious play converge to a set which can be characterized in terms of the ϵ -equilibrium set of the game, and the level sets of the potential function of a close potential game. We also establish that for games sufficiently close to a potential game, the empirical frequencies of fictitious play converge to a neighborhood of the (mixed) equilibrium set. Moreover, the size of this neighborhood depends on the distance of the original game from a nearby potential game. This generalizes the result of Monderer and Shapley (1996a), on convergence of empirical frequencies to mixed Nash equilibria in potential games.

In this chapter, we only consider the discrete-time version of fictitious play, i.e., the update process starts at a given strategy profile at time $t = 0$, and players can update their strategies at discrete time instants $t \in \{1, 2, \dots\}$. Throughout this subsection we denote the strategy used by player m at time instant t by p_t^m , and we denote by $\mathbf{1}(p_t^m = p^m)$ the indicator function which equals to 1 if $p_t^m = p^m$, and 0 otherwise. A formal definition of discrete-time fictitious play dynamics is given next.

Definition 3.5.3 (Discrete-Time Fictitious Play). Let $\mu_T^m(q^m) = \frac{1}{T} \sum_{t=0}^{T-1} \mathbf{1}(p_t^m = q^m)$ denote the empirical frequency that player m uses strategy q^m from time instant 0 to time instant $T - 1$, and μ_T^{-m} denote the collection of empirical frequencies of all players but m . A game play, where at each time instant t , every player m , chooses a strategy p_t^m such that

$$p_t^m \in \arg \max_{q^m \in E^m} u^m(q^m, \mu_t^{-m})$$

is referred to as *discrete-time fictitious play*. That is, *fictitious play dynamics* is the update process, where each player chooses its best response to the empirical frequencies of the actions of other players.

We refer to μ_t^m as the distribution of empirical frequencies of player m 's strategies at time t . Note that μ_t^m can be thought of as vector with length $|E^m|$, whose entries are indexed by strategies of player m , i.e., $\mu_t^m(p^m)$ denotes the entry of the vector corresponding to the empirical frequency player m uses strategy p^m with. Similarly, we define the joint empirical frequency distribution of all players as $\mu_t = \{\mu_t^m\}_{m \in \mathcal{M}}$. Note that $\mu_t^m \in \Delta E^m$, i.e., empirical frequency distributions are mixed strategies, and similarly $\mu_t \in \prod_{m \in \mathcal{M}} \Delta E^m$.

Observe that the evolution of this empirical frequency distribution can be captured by the following equation:

$$\mu_{t+1} = \frac{t}{t+1} \mu_t + \frac{1}{t+1} I_t, \quad (3.9)$$

where $I_t = \{I_t^m\}_{m \in \mathcal{M}}$, and I_t^m is a vector which has the same size as μ_t^m and its entry corresponding to strategy p^m is given by $I_t^m(p^m) = \mathbf{1}(p_t^m = p^m)$. Rearranging the terms in (3.9), and observing that $I_t, \mu_t \in \prod_{m \in \mathcal{M}} \Delta E^m$ are vectors with entries in $[0, 1]$ we conclude

$$\|\mu_{t+1} - \mu_t\| = \frac{1}{t+1} \|I_t - \mu_t\| = O\left(\frac{1}{t}\right), \quad (3.10)$$

where $O(\cdot)$ stands for the big-O notation, i.e., $f(x) = O(g(x))$, implies that there exists some x_0 and a constant c such that $|f(x)| \leq c|g(x)|$ for all $x \geq x_0$.

We start analyzing discrete-time fictitious play in near-potential games, by first focusing on the change in the value of the potential function along the fictitious play updates in the original game. In particular, we show that in near-potential games if the empirical frequencies are outside some ϵ -equilibrium set, then the potential of the close potential game (evaluated at the empirical frequency distribution) increases by discrete-time fictitious play updates.¹²

Lemma 3.5.3. *Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a close potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the potential function of $\hat{\mathcal{G}}$ by ϕ . Assume that in \mathcal{G} players update their strategies according to discrete-time fictitious play dynamics, and at some time instant $T > 0$, the empirical frequency distribution μ_T is outside an ϵ -equilibrium set of \mathcal{G} . Then,*

$$\phi(\mu_{T+1}) - \phi(\mu_T) \geq \frac{\epsilon - M\delta}{T+1} + O\left(\frac{1}{T^2}\right).$$

Proof. Consider the mixed extension of the potential function $\phi(\mathbf{x}) = \sum_{\mathbf{p} \in E} \phi(\mathbf{p}) \prod_{m \in \mathcal{M}} x^m(p^m)$, where $\mathbf{x} = \{x^m\}_m$ and $x^m(p^m)$ denotes the probability player m plays strategy p^m . The

¹²Our approach here is similar to the one used in Monderer and Shapley (1996a) to analyze discrete-time fictitious play in potential games.

expression for $\phi(\mathbf{x})$ implies that Taylor expansion of ϕ around μ_T satisfies

$$\phi(\mu_{T+1}) = \phi(\mu_T) + \sum_{m \in \mathcal{M}} \sum_{p^m \in E^m} (\mu_{T+1}^m(p^m) - \mu_T^m(p^m)) \phi(p^m, \mu_T^{-m}) + O(\|\mu_{T+1} - \mu_T\|^2).$$

Observing from (3.9) that $\mu_{t+1} - \mu_t = \frac{1}{t+1}(I_t - \mu_t)$, and noting from (3.10) that $\|\mu_{t+1} - \mu_t\| = O(\frac{1}{t})$ the above equality can be rewritten as

$$\phi(\mu_{T+1}) = \phi(\mu_T) + \sum_{m \in \mathcal{M}} \sum_{p^m \in E^m} \frac{1}{T+1} (\mathbf{1}(p_T^m = p^m) - \mu_T^m(p^m)) \phi(p^m, \mu_T^{-m}) + O\left(\frac{1}{T^2}\right).$$

Rearranging the terms, and noting that $\sum_{p^m \in E^m} \mu_T^m(p^m) \phi(p^m, \mu_T^{-m}) = \phi(\mu_T^m, \mu_T^{-m})$, it follows that

$$\begin{aligned} \phi(\mu_{T+1}) &= \phi(\mu_T) + \sum_{m \in \mathcal{M}} \frac{1}{T+1} \phi(p_T^m, \mu_T^{-m}) - \sum_{m \in \mathcal{M}} \frac{1}{T+1} \phi(\mu_T^m, \mu_T^{-m}) + O\left(\frac{1}{T^2}\right) \\ &= \phi(\mu_T) + \frac{1}{T+1} \sum_{m \in \mathcal{M}} (\phi(p_T^m, \mu_T^{-m}) - \phi(\mu_T^m, \mu_T^{-m})) + O\left(\frac{1}{T^2}\right). \end{aligned}$$

Since $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$, the above equality and the definition of MPD imply

$$\phi(\mu_{T+1}) \geq \phi(\mu_T) + \frac{1}{T+1} \sum_{m \in \mathcal{M}} (u^m(p_T^m, \mu_T^{-m}) - u^m(\mu_T^m, \mu_T^{-m}) - \delta) + O\left(\frac{1}{T^2}\right). \quad (3.11)$$

By definition of the fictitious play dynamics, every player m plays its best response to μ_T^{-m} , therefore $u^m(p_T^m, \mu_T^{-m}) - u^m(\mu_T^m, \mu_T^{-m}) \geq 0$ for all m . Additionally, if μ_T is outside the ϵ -equilibrium set, as in the statement of the lemma, then it follows that $u^m(p_T^m, \mu_T^{-m}) - u^m(\mu_T^m, \mu_T^{-m}) \geq \epsilon$ for at least one player. Therefore, (3.11) implies

$$\phi(\mu_{T+1}) \geq \phi(\mu_T) + \frac{\epsilon - M\delta}{T+1} + O\left(\frac{1}{T^2}\right),$$

hence, the claim follows. \square

The above theorem implies that if μ_T is not in the ϵ -equilibrium set for some $\epsilon > M\delta$, and T sufficiently large, then the potential evaluated at empirical frequencies increases when players update their strategies. Since the mixed extension of the potential is a bounded function, the potential cannot increase unboundedly, and this observation suggests that the ϵ -equilibrium set is eventually reached by the empirical frequency distribution. On the other hand, at a later time instant μ_T can still leave this equilibrium set, and before it does so the potential cannot be lower than the lowest potential in this set (since μ_T itself belongs to this set). Moreover, after μ_T leaves the ϵ -equilibrium set the potential keeps increasing. Thus, the empirical frequencies are contained in the set of mixed strategy profiles, which have

potential at least as large as the minimum potential in this approximate equilibrium set. We next make this intuition precise, and characterize the set of limiting mixed strategies for fictitious play in near-potential games. We adopt the following convergence notion: we say that empirical frequencies of fictitious play converge to a set $S \subset \prod_{m \in \mathcal{M}} \Delta E^m$, if $\inf_{\mathbf{x} \in S} \|\mu_t - \mathbf{x}\| \rightarrow 0$ as $t \rightarrow \infty$.

Theorem 3.5.1. *Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a close potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the potential function of $\hat{\mathcal{G}}$ by ϕ . Assume that in \mathcal{G} players update their strategies according to discrete-time fictitious play dynamics, and let \mathcal{X}_α denote the α -equilibrium set of \mathcal{G} . For any $\epsilon > 0$, there exists a time instant $T_\epsilon > 0$ such that for all $t > T_\epsilon$*

$$\mu_t \in C_\epsilon \triangleq \left\{ \mathbf{x} \in \prod_{m \in \mathcal{M}} \Delta E^m \mid \phi(\mathbf{x}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta + \epsilon}} \phi(\mathbf{y}) \right\}.$$

Proof. Let ϵ' be such that $\epsilon > \epsilon' > 0$. It can be seen from the definition of C_ϵ that $\mathcal{X}_{M\delta + \epsilon'} \subset \mathcal{X}_{M\delta + \epsilon} \subset C_\epsilon$. We prove the claim in two steps: (i) We first show that in this update process $\mathcal{X}_{M\delta + \epsilon'}$ is visited infinitely often by μ_t , i.e., for all T' , there exists $t > T'$ such that $\mu_t \in \mathcal{X}_{M\delta + \epsilon'}$, (ii) We prove that there exists a T'' such that if $\mu_t \in C_\epsilon$ for some $t > T''$, then for all $t' > t$ we have $\mu_{t'} \in C_\epsilon$. Thus, the second step guarantees that if C_ϵ is visited at a sufficiently later time instant, then μ_t remains in C_ϵ . Since $\mathcal{X}_{M\delta + \epsilon'} \subset C_\epsilon$ the first step ensures that such a time instant exists, and the claim in the theorem immediately follows from (ii). Moreover, this time instant corresponds to T_ϵ in the theorem statement.

Proof of both steps rely on the following simple observation: Lemma 3.5.3 implies that there exists a large enough T , such that if the empirical frequencies do not belong to $\mathcal{X}_{M\delta + \epsilon'}$ at a time instant $t > T$, then ϕ increases:

$$\phi(\mu_{t+1}) - \phi(\mu_t) \geq \frac{M\delta + \epsilon' - M\delta}{(t+1)} + O\left(\frac{1}{t^2}\right) > \frac{\epsilon'}{2(t+1)} > 0. \quad (3.12)$$

We prove (i) by contradiction. Assume that there exists a T' such that $\mu_t \notin \mathcal{X}_{M\delta + \epsilon'}$ for $t > T'$, and let $T_m = \max\{T, T'\}$. Then, (3.12) holds for all $t = \{T_m + 1, \dots\}$, and summing both sides of this inequality over this set we obtain

$$\limsup_{t \rightarrow \infty} \phi(\mu_{t+1}) - \phi(\mu_{T_m+1}) \geq \sum_{t=T_m+1}^{\infty} \frac{\epsilon'}{2(t+1)}.$$

Since the mixed extension of the potential is a bounded function, it follows that the left hand side of the above inequality is bounded, but the right hand side grows unboundedly. Hence, we reach a contradiction, and (i) follows.

Lemma 3.5.1 (ii) implies that there exists some $\theta > 0$ such that if a strategy profile \mathbf{x} is an $(M\delta + \epsilon')$ -equilibrium, then any strategy profile \mathbf{y} that satisfies $\|\mathbf{x} - \mathbf{y}\| < \theta$ is an $(M\delta + \epsilon)$ -equilibrium (recall that $\epsilon > \epsilon' > 0$). Since $\|\mu_{t+1} - \mu_t\| = O(1/t)$ by (3.10), this

implies that there exists some $T'' > T$, such that for all $t > T''$ if $\mu_t \in \mathcal{X}_{M\delta+\epsilon'}$, then we have

$$\mu_{t+1} \in \mathcal{X}_{M\delta+\epsilon}. \quad (3.13)$$

Let $\mu_t \in C_\epsilon$ for some time instant $t > T''$. If $\mu_t \in \mathcal{X}_{M\delta+\epsilon'}$, then by (3.13) $\mu_{t+1} \in \mathcal{X}_{M\delta+\epsilon} \subset C_\epsilon$. If, on the other hand, $\mu_t \in C_\epsilon - \mathcal{X}_{M\delta+\epsilon'}$, then by (3.12) and the definition of C_ϵ we have

$$\phi(\mu_{t+1}) > \phi(\mu_t) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta+\epsilon}} \phi(\mathbf{y}), \quad (3.14)$$

and hence $\mu_{t+1} \in C_\epsilon$. Thus, we have established that there exists some T'' such that if $\mu_t \in C_\epsilon$ for some $t > T''$, then $\mu_{t+1} \in C_\epsilon$, and hence (ii) follows. \square

The above theorem establishes that after finite time μ_t is contained in the set C_ϵ for any $\epsilon > 0$. Corollary 3.5.1, establishes that in the limit this result can be strengthened: as $t \rightarrow \infty$, μ_t converges to a set, which is a subset of C_ϵ for every $\epsilon > 0$. The proof can be found in Section 3.7.

Corollary 3.5.1. *The empirical frequencies of discrete-time fictitious play converge to*

$$C \triangleq \left\{ \mathbf{x} \in \prod_{m \in \mathcal{M}} \Delta E^m \mid \phi(\mathbf{x}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) \right\}.$$

This result suggests that in near-potential games, the empirical frequencies of fictitious play converge to a set where the potential is at least as large as the minimum potential in an approximate equilibrium set. For exact potential games, it is known that the empirical frequencies converge to a Nash equilibrium (Monderer and Shapley, 1996a). It can be seen from Definition 2.2.1 that in potential games, maximizers of the potential function are equilibria of the game. Thus, in potential games with a unique equilibrium the equilibrium is the unique maximizer of the potential function. Hence, for such games, we have $\delta = 0$, $\min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) = \max_{\mathbf{x} \in \prod_{m \in \mathcal{M}} \Delta E^m} \phi(\mathbf{x})$, and Corollary 3.5.1 implies that empirical frequencies of fictitious play converge to the unique equilibrium of the game, recovering the convergence result of Monderer and Shapley (1996a). However, when there are multiple equilibria Corollary 3.5.1 suggests that empirical frequencies converge to the set of mixed strategy profiles that have potential weakly larger than the minimum potential attained by the equilibria. While this set contains equilibria, it may contain a continuum of other mixed strategy profiles. This suggests that in games with multiple equilibria our result may provide a loose characterization of the limiting behavior of fictitious play dynamics.

We next show that by exploiting the properties of mixed approximate equilibrium sets, it is possible to obtain a stronger result. Before we present our result, we discuss a feature of mixed equilibrium sets which will be key in our analysis: For small ϵ , the ϵ -equilibrium set is contained in a small neighborhood of equilibria (this statement follows from Lemma 3.5.2 (ii) by considering the upper semicontinuity of the approximate equilibrium correspondence

$g(\alpha)$ at $\alpha = 0$). This property is illustrated in Example 3.5.1.

Example 3.5.1 (Mixed equilibrium set of Battle of the Sexes:). *Consider the two-player battle of the sexes (BoS) game: Each player has two possible actions $\{O, F\}$, and the payoffs of players are as given in Table 3.3. This game has three equilibria: (i) both players use*

	O	F
O	3, 2	0, 0
F	0, 0	2, 3

Table 3.3: Payoffs in BoS.

O, (ii) Both players use F, (iii) Row player uses O with probability 0.6, and column player uses O with probability 0.4. Note that since this is a game where each player has only two strategies, the probability of using strategy O, in the third case uniquely identifies the corresponding mixed equilibrium. For different values of ϵ , the set of ϵ -equilibria of this game is shown in Figure 3-4. It follows that the set of ϵ -equilibria is contained in disjoint neighborhoods of equilibria for small values of ϵ .

It was established in Lemma 3.5.3 that the potential function of a nearby potential game (with MPD δ to the original game), evaluated at the empirical frequency distribution, increases when this distribution is outside the $M\delta$ -equilibrium set of the original game (where M is the number of players). If δ is sufficiently small, then the $M\delta$ -equilibria of the game will be contained in a small neighborhood of the equilibria, as illustrated above and shown in Lemma 3.5.2 (ii). Thus, for sufficiently small δ , it is possible to establish that the potential of a close potential game increases outside a small neighborhood of the equilibria of the game. In Theorem 3.5.2, we use this observation to show that for sufficiently small δ the empirical frequencies of fictitious play dynamics converge to a neighborhood of an equilibrium. We state the theorem under the assumption that the original game has finitely many equilibria. This assumption generically holds, i.e., for any game a (nondegenerate) random perturbation of payoffs will lead to such a game with probability one (see Fudenberg and Tirole (1991)).

When stating our result, we make use of the Lipschitz continuity of the mixed extension of the potential function, as established in Lemma 3.5.1. We also make use of a function $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, which quantifies the size of the neighborhood of equilibria which contains the approximate equilibrium sets of games. For a game \mathcal{G} with l equilibria, denoted by $\mathbf{x}_1, \dots, \mathbf{x}_l$, this function can be formally defined as follows:

$$f(\alpha) = \max_{\mathbf{x} \in \mathcal{X}_\alpha} \min_{k \in \{1, \dots, l\}} \|\mathbf{x} - \mathbf{x}_k\|, \quad (3.15)$$

for all $\alpha \in \mathbb{R}_+$. Note that $\min_{k \in \{1, \dots, l\}} \|\mathbf{x} - \mathbf{x}_k\|$ is continuous in \mathbf{x} , since it is minimum of finitely many continuous functions. Moreover, \mathcal{X}_α is a compact set, since ϵ -equilibria

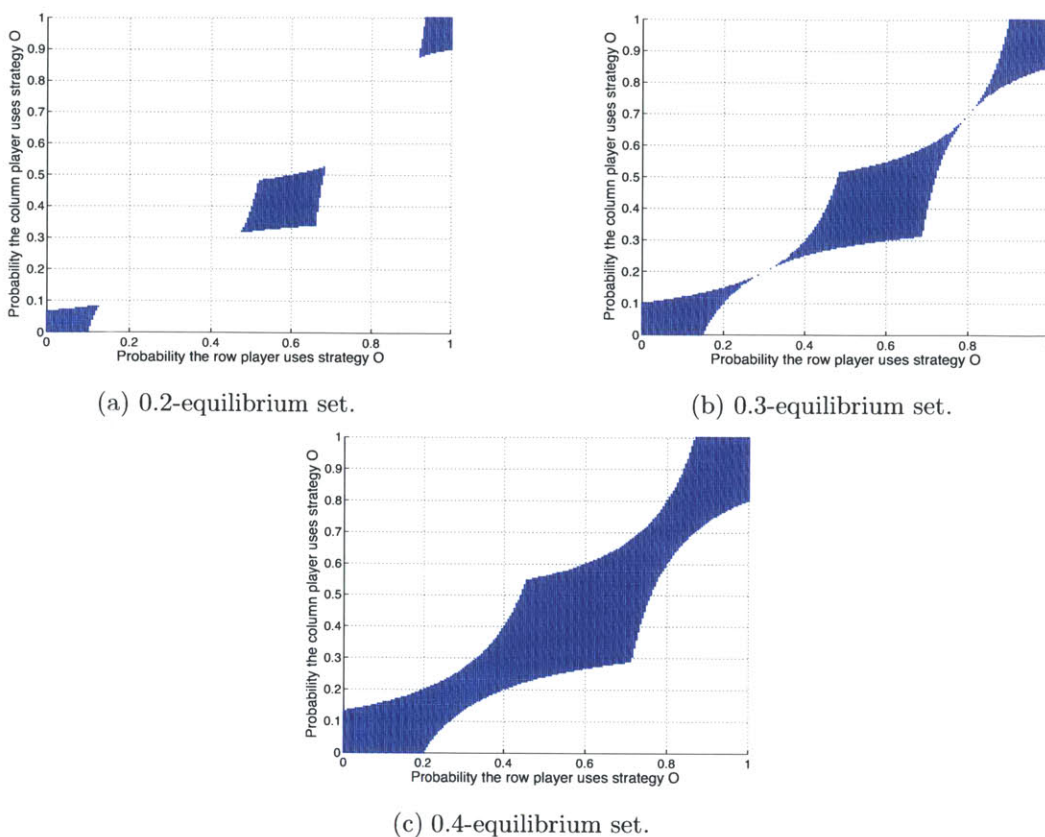


Figure 3-4: Approximate equilibrium sets in BoS are contained in disjoint neighborhoods of equilibria for small ϵ .

are defined by finitely many inequality constraints of the form (3.8). Therefore, in (3.15) maximum is achieved and f is well-defined for all $\alpha \geq 0$.

Additionally, we define two variables, (a, d) , which characterize the approximate equilibrium sets of the underlying game \mathcal{G} : (i) the minimum pairwise distance between the equilibria is denoted by $d \triangleq \min_{i \neq j} \|\mathbf{x}_i - \mathbf{x}_j\|$ (ii) $a \triangleq \sup\{\alpha | f(\alpha) < d/4\} > 0$, i.e., for every $\alpha < a$, the α -equilibrium is at most $d/4$ distant from an equilibrium of \mathcal{G} . Next, using these definitions, we state an improved convergence result for fictitious play in near-potential games.

Theorem 3.5.2. *Consider a game \mathcal{G} and let $\hat{\mathcal{G}}$ be a close potential game such that $d(\mathcal{G}, \hat{\mathcal{G}}) \leq \delta$. Denote the potential function of $\hat{\mathcal{G}}$ by ϕ , and the Lipschitz constant of the mixed extension of ϕ by L . Assume that \mathcal{G} has finitely many equilibria, and in \mathcal{G} players update their strategies according to discrete-time fictitious play dynamics.*

(i) *There exists some $\bar{\delta} > 0$, and $\bar{\epsilon} > 0$ satisfying*

$$M\bar{\delta} + \bar{\epsilon} < a, \quad \text{and} \quad f(M\bar{\delta} + \bar{\epsilon}) < \frac{(a - M\bar{\delta})d}{24LM}.$$

(ii) Consider any $\bar{\delta} > 0$, and $\bar{\epsilon} > 0$ satisfying (i). Provided that $\bar{\delta} \geq \delta \geq 0$, it can be established that the empirical frequencies of fictitious play converge to

$$\left\{ \mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_k\| \leq \frac{4f(M\delta)ML}{\epsilon} + f(M\delta + \epsilon), \text{ for some equilibrium } \mathbf{x}_k \right\}, \quad (3.16)$$

for any ϵ , such that $\bar{\epsilon} \geq \epsilon > 0$.

The proof of this theorem can be found in Section 3.7.¹³

The proof has three main steps illustrated in Figures 3-5 and 3-6. As explained earlier, for small δ and ϵ , the $M\delta + \epsilon$ -equilibrium set of the game is contained in disjoint neighborhoods of the equilibria of the game. Lemma 3.5.3 implies that potential evaluated at μ_t increases outside this approximate equilibrium set with strategy updates. In the proof, we first quantify the increase in the potential, when μ_t leaves this approximate equilibrium set and returns back to it at a later time instant (see Figure 3-5a). Then, using this increase condition we show that for sufficiently large t , μ_t can visit the approximate equilibrium set infinitely often only around one equilibrium, say $\mathbf{x}_{k'}$ (see Figure 3-5b). This holds since, the increase condition guarantees that the potential increases significantly when μ_t leaves the neighborhood of an equilibrium \mathbf{x}_k , and reaches to that of $\mathbf{x}_{k'}$. Finally, using the increase condition one more time, we establish that if after time T , μ_t visits the approximate equilibrium set only in the neighborhood of $\mathbf{x}_{k'}$, we can construct a neighborhood of $\mathbf{x}_{k'}$, which contains μ_t for all $t > T$ (see Figure 3-6). In equation (3.16) of Theorem 3.5.2, we provide bounds on this neighborhood, as a function of δ (that characterizes the ‘‘closeness’’ of the original game to a potential game), and f (that captures how the size of the ϵ -equilibrium sets increase, as a function of ϵ).

Observe that if $\delta = 0$, i.e., the original game is a potential game, then $f(M\delta) = 0$, and Theorem 3.5.2 implies that empirical frequencies of fictitious play converge to the $f(\epsilon)$ -neighborhood of equilibria for any ϵ such that $\bar{\epsilon} \geq \epsilon > 0$. Thus, choosing ϵ arbitrarily small, and observing that $\lim_{x \rightarrow 0} f(x) = 0$, our result implies that in potential games, empirical frequencies converge to the set of Nash equilibria. Hence, as a special case of Theorem 3.5.2, we obtain the convergence result of Monderer and Shapley (1996a).

Assume that $\delta \neq 0$ and a small $\epsilon < \bar{\epsilon}$ is given. If δ is sufficiently small then $f(M\delta)/\epsilon \approx 0$, since $\lim_{x \rightarrow 0} f(x) = 0$. Consequently, $\frac{4f(M\delta)ML}{\epsilon} + f(M\delta + \epsilon)$ is small, and Theorem 3.5.2 establishes convergence of empirical frequencies to a small neighborhood of equilibria. Thus, we conclude that for games that are close to potential games, i.e., for $\delta \ll 1$, Theorem 3.5.2 establishes convergence of empirical frequencies to a small neighborhood of equilibria.

Corollary 3.5.1 and Theorem 3.5.2 give a systematic framework for approximately characterizing the limiting behavior of fictitious play in arbitrary games. Moreover, such a

¹³A strand of the literature characterizes the limiting behavior of discrete time fictitious play by exploiting its relation to a continuous time update rule (see for instance Benaïm et al. (2005)). In our proof, we instead follow a direct approach, which exploits Lemma 3.5.3, and provides a quantitative characterization of the limiting behavior of fictitious play dynamics for near-potential games.

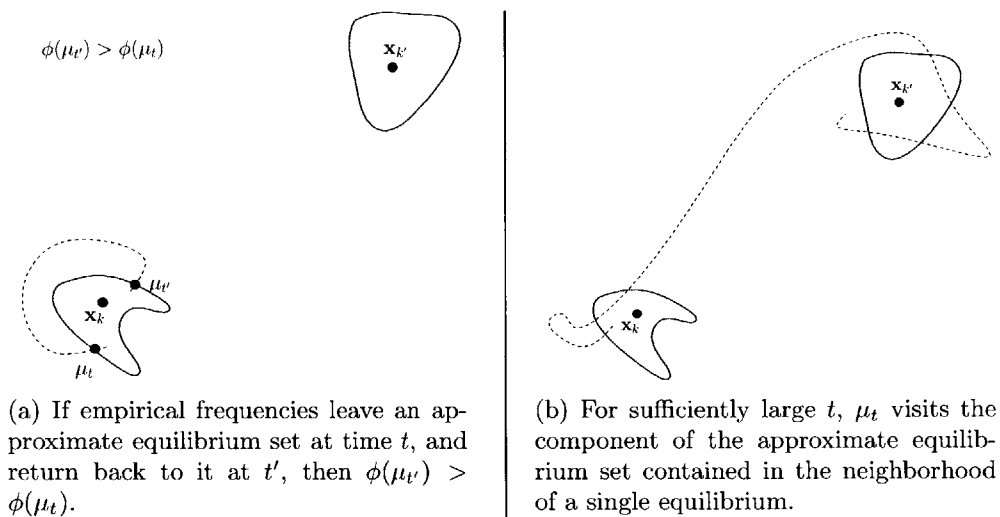


Figure 3-5: For small δ and ϵ , $M\delta + \epsilon$ -equilibrium set (enclosed by solid lines around equilibria $\mathbf{x}_{k'}$ and \mathbf{x}_k) is contained in disjoint neighborhoods of equilibria. If the empirical frequency distribution, μ_t , is outside this approximate equilibrium set, then the potential increases with each strategy update. Assume that empirical frequency distribution leaves an approximate equilibrium set (at time t) and returns back to it at a later time instant ($t' > t$). We first quantify the resulting increase in the potential (left). If μ_t travels from the component of the approximate equilibrium set in the neighborhood of equilibrium \mathbf{x}_k to that in the neighborhood of equilibrium $\mathbf{x}_{k'}$, then the increase in the potential is significant, and consequently μ_t cannot visit the approximate equilibrium set in the neighborhood of equilibrium \mathbf{x}_k at a later time instant (right).




Figure 3-6: If after time T , μ_t only visits the approximate equilibrium set in the neighborhood of a single equilibrium $\mathbf{x}_{k'}$, then we can establish that μ_t never leaves a neighborhood of this equilibrium for $t > T$. The size of this neighborhood is denoted by r in the figure and can be expressed as in Theorem 3.5.2.

characterization can be obtained even in settings where the underlying game does not share similar ordinal properties to potential games. Following a similar argument as in the case of logit response dynamics, our result also allows for characterizing robustness of convergence results for potential games to payoff perturbations.

3.6 Summary

In this chapter, we presented a framework for studying the limiting behavior of adaptive learning dynamics in finite strategic form games by exploiting their relation to nearby potential games. We restricted our attention to better/best response, logit response and fictitious play dynamics. We showed that for near-potential games trajectories of better/best response dynamics converge to ϵ -equilibrium sets, where ϵ depends on closeness to a potential game. We studied the stochastically stable strategy profiles of logit response dynamics and proved that they are contained in the set of strategy profiles that approximately maximize the potential function of a nearby potential game. In the case of fictitious play we focused on the empirical frequencies of players' actions, and established that they converge to a small neighborhood of equilibria in near-potential games. Our results suggest that games that are close to a potential game inherit the dynamical properties (such as convergence to approximate equilibrium sets) of potential games. Additionally, since a close potential game to a given game can be found by using the decomposition approach of Chapter 2, or solving a convex optimization problem, as discussed in Section 3.2, this enables us to characterize the dynamical properties of strategic form games by first identifying a nearby potential game to this game, and then studying the dynamical properties of the nearby potential game.

3.7 Appendix: Proofs of Section 3.5

Proof of Corollary 3.5.1: Let $\epsilon_n = M\delta + \frac{1}{n}$ for $n \in \mathbb{Z}_+$. Observe that since the mixed extension of the potential function is continuous, C and C_{ϵ_n} are closed sets for any $n \in \mathbb{Z}_+$. Since C is closed $\min_{\mathbf{y} \in C} \|\mathbf{x} - \mathbf{y}\|$ is well-defined for any $\mathbf{x} \in \prod_{m \in \mathcal{M}} \Delta E^m$.

We claim that for any $\theta > 0$ the set

$$S_\theta = \left\{ \mathbf{x} \in \prod_{m \in \mathcal{M}} \Delta E^m \mid \min_{\mathbf{y} \in C} \|\mathbf{x} - \mathbf{y}\| < \theta \right\}, \quad (3.17)$$

is such that $C_{\epsilon_n} \subset S_\theta$ for some n . Note that if this claim holds, then it follows from Theorem 3.5.1 that there exists some T_θ such that for all $t > T_\theta$ we have $\mu_t \in S_\theta$. Using the definition of S_θ given in (3.17), this implies

$$\limsup_{t \rightarrow \infty} \min_{\mathbf{x} \in C} \|\mathbf{x} - \mu_t\| < \theta. \quad (3.18)$$

Moreover, since $\theta > 0$ is arbitrary, and $\|\mathbf{x} - \mu_t\| \geq 0$, using (3.18) we obtain

$$\lim_{t \rightarrow \infty} \min_{\mathbf{x} \in C} \|\mathbf{x} - \mu_t\| = 0.$$

Thus, if we prove $C_{\epsilon_n} \subset S_\theta$ for some n , it follows that μ_t converges to C .

In order to prove $C_{\epsilon_n} \subset S_\theta$ we first obtain a certificate which can be used to guarantee that a mixed strategy profile belongs to S_θ . Then, we show that for large enough n any $\mathbf{z} \in C_{\epsilon_n}$ satisfies this certificate, and hence belongs to S_θ .

It follows from Lemma 3.5.2 (i) (by setting $\nu = \phi$ and $v = -\min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y})$) and definition of upper semicontinuity (Definition 3.5.2) that there exists $\gamma > 0$ such that θ neighborhood of $\{\mathbf{x} | \phi(\mathbf{x}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y})\}$ contains $\{\mathbf{x} | \phi(\mathbf{x}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma\}$. Hence, for any \mathbf{z} satisfying $\phi(\mathbf{z}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma$ there exists some \mathbf{x} satisfying $\phi(\mathbf{x}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y})$ and $\|\mathbf{x} - \mathbf{z}\| < \theta$. Note that the definition of S_θ implies that \mathbf{z} for which there exists such \mathbf{x} belongs to S_θ . Thus, if $\phi(\mathbf{z}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma$ it follows that $\mathbf{z} \in S_\theta$.

We next show that for large enough n , any \mathbf{z} which belongs to C_{ϵ_n} , satisfies the above certificate and hence belongs to S_θ . Let L denote the Lipschitz constant for the mixed extension of ϕ , as given in Lemma 3.5.1 (i), and define $\theta' = \gamma/L > 0$. Lemma 3.5.2 (ii) and Definition 3.5.2 imply that for large enough n , $\mathcal{X}_{M\delta + \frac{1}{n}}$ is contained in θ' neighborhood of $\mathcal{X}_{M\delta}$, i.e., if $\mathbf{y} \in \mathcal{X}_{M\delta + \frac{1}{n}}$ then there exists $\mathbf{x} \in \mathcal{X}_{M\delta}$ such that $\|\mathbf{x} - \mathbf{y}\| < \theta'$. Moreover, by Lemma 3.5.1 (i), it follows that $\phi(\mathbf{y}) \geq \phi(\mathbf{x}) - L\theta' = \phi(\mathbf{x}) - \gamma$. Thus, we conclude that there exists large enough n such that

$$\min_{\mathbf{y} \in \mathcal{X}_{M\delta + 1/n}} \phi(\mathbf{y}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma. \quad (3.19)$$

Let $\mathbf{z} \in C_{\epsilon_n}$ for some n for which (3.19) holds. By definition of C_ϵ it follows that $\phi(\mathbf{z}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta + 1/n}} \phi(\mathbf{y})$. Thus, (3.19) implies that $\phi(\mathbf{z}) \geq \min_{\mathbf{y} \in \mathcal{X}_{M\delta}} \phi(\mathbf{y}) - \gamma$. However, as argued before such \mathbf{z} belong to S_θ . Hence, we have established that for large enough n , if $\mathbf{z} \in C_{\epsilon_n}$ then $\mathbf{z} \in S_\theta$. Therefore, the claim follows. \square

Proof of Theorem 3.5.2: From the definition of f , it follows that the union of closed balls of radius $f(\alpha)$, centered at equilibria, contain α -equilibrium set of the game. Thus, intuitively, $f(\alpha)$ captures the size of a closed neighborhood of equilibria, which contains α -equilibria of the underlying game. This is illustrated in Figure 3-7.

As stated in the theorem statement, we define the minimum pairwise distance between the equilibria as $d \triangleq \min_{i \neq j} \|\mathbf{x}_i - \mathbf{x}_j\|$, and $a = \sup\{\alpha | f(\alpha) < d/4\}$. Lemma 3.5.2 (ii) implies (using upper semicontinuity at 0) that $\alpha > 0$ such that $f(\alpha) < d/4$ exists and hence $a > 0$. Since d is defined as the minimum pairwise distance between the equilibria, it follows that α -equilibria of the game are contained in disjoint $f(\alpha) < d/4$ neighborhoods around equilibria of the game (for $\alpha < a$), i.e., if $\mathbf{x} \in \mathcal{X}_\alpha$, then $\|\mathbf{x} - \mathbf{x}_k\| \leq f(\alpha)$ for exactly one equilibrium \mathbf{x}_k . Moreover, for $\alpha_1 \leq \alpha$, since $\mathcal{X}_{\alpha_1} \subset \mathcal{X}_\alpha$, it follows that α_1 -equilibria of the




Figure 3-7: Consider a game with a unique equilibrium \mathbf{x}_k . The α -equilibrium set of the game (enclosed by a solid line around \mathbf{x}_k) is contained in the $f(\alpha)$ neighborhood of this equilibrium.

game are contained in disjoint neighborhoods of equilibria.

We prove the theorem in 5 steps summarized below. First two steps explore the properties of function f , and establish existence of $\bar{\delta}$ and $\bar{\epsilon}$ presented in the theorem statement. Last three steps are the main steps of the proof, where we establish convergence of fictitious play to a neighborhood of equilibria.

- *Step 1:* We first show that f is (i) weakly increasing, (ii) upper semicontinuous, and it satisfies (iii) $f(0) = 0$, (iv) $f(x) \rightarrow 0$ as $x \rightarrow 0$.
- *Step 2:* We show that there exists some $\bar{\delta} > 0$ and $\bar{\epsilon} > 0$ such that the following inequalities hold:

$$M\bar{\delta} + \bar{\epsilon} < a, \quad (3.20)$$

and

$$f(M\bar{\delta} + \bar{\epsilon}) < \frac{(a - M\bar{\delta})d}{24LM}. \quad (3.21)$$

We will prove the statement of the theorem assuming that $0 \leq \delta < \bar{\delta}$, and establish convergence to the set in (3.16), for any ϵ such that $0 < \epsilon \leq \bar{\epsilon}$. As can be seen from the definition of a and f (see (3.15)), the first inequality guarantees that $M\bar{\delta} + \bar{\epsilon}$ -equilibrium set is contained in disjoint neighborhoods of equilibria, and the second one guarantees that these neighborhoods are small. In Step 4, we will exploit this observation, and use the inequalities in (3.20) and (3.21) to establish that the empirical frequency distribution μ_t can visit the component of $\mathcal{X}_{M\bar{\delta} + \bar{\epsilon}}$ contained in the neighborhood of only a single equilibrium infinitely often.

- *Step 3:* Let ϵ_1, ϵ_2 be such that $\epsilon_2 > \epsilon_1 > 0$. Assume that (i) at some time instant T , μ_t is contained in $\mathcal{X}_{M\delta + \epsilon_1}$, (ii) at time instants T_1 and T_2 (such that $T_2 > T_1 > T$) μ_t leaves $\mathcal{X}_{M\delta + \epsilon_1}$ and $\mathcal{X}_{M\delta + \epsilon_2}$ respectively and (iii) at time instants T'_2 and T'_1 (such that

$T'_1 > T'_2 > T_2$) μ_t returns back to $\mathcal{X}_{M\delta+\epsilon_2}$ and $\mathcal{X}_{M\delta+\epsilon_1}$ respectively. In Figure 3-8, the path μ_t follows between T_1 and T'_1 is illustrated.

In this step, we provide a lower bound on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$, i.e., the increase in the potential when μ_t follows such a path. This lower bound holds for any ϵ_1 and ϵ_2 provided that $\epsilon_2 > \epsilon_1 > 0$. We use this result by choosing different values for ϵ_1 and ϵ_2 in Steps 4 and 5.

Our lower bound in Step 3 is a function of ϵ_2 . In addition to this lower bound, in Steps 4 and 5, we use the $M\delta + \epsilon_1$ equilibrium set and Lipschitz continuity of the potential to provide an upper bound on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$ as a function of ϵ_1 . Thus, properties of $M\delta + \epsilon_1$ and $M\delta + \epsilon_2$ equilibrium sets are exploited for obtaining upper and lower bounds on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$ respectively. We establish convergence of fictitious play updates to a neighborhood of an equilibrium by using these bounds together in Steps 4 and 5. We emphasize that allowing for two different approximate equilibrium sets leads to better bounds on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$, and a more informative characterization of the limiting behavior of fictitious play, as opposed to using a single approximate equilibrium set, i.e., setting $\epsilon_1 = \epsilon_2$.

- *Step 4:* Our objective in this step is to establish that fictitious play can visit the component of an approximate equilibrium set contained in the neighborhood of only one equilibrium infinitely often.

Let $\epsilon_1 = \bar{\epsilon}$ and $\epsilon_2 = a - M\bar{\delta}$. By (3.20) we have $\epsilon_1 < \epsilon_2$, and using the definition of a we establish that $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ are contained in disjoint neighborhoods of equilibria. Assume that μ_t leaves the components of $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ in the neighborhood of equilibrium \mathbf{x}_k , and reaches to a similar neighborhood around equilibrium $\mathbf{x}_{k'}$. Using Step 3 we establish a lower bound on the increase in the potential when μ_t follows such a trajectory. We also provide an upper bound, using the Lipschitz continuity of the potential and inequalities (3.20) and (3.21). Comparing these bounds, we establish that the maximum potential in the neighborhood of equilibrium \mathbf{x}_k is lower than the minimum potential in the neighborhood of $\mathbf{x}_{k'}$. Since, \mathbf{x}_k and $\mathbf{x}_{k'}$ are arbitrary, this observation implies that μ_t cannot visit the component of $\mathcal{X}_{M\delta+\epsilon_1}$ contained in the neighborhood of \mathbf{x}_k at a later time instant. Hence, it follows that μ_t visits only one such component infinitely often.

- *Step 5:* In this step we show that μ_t converges to the approximate equilibrium set given in the theorem statement.

Let ϵ_1, ϵ_2 be such that $0 < \epsilon_1 < \epsilon_2 \leq \bar{\epsilon}$. We consider the equilibrium, whose neighborhood is visited infinitely often (as obtained in Step 4), and a trajectory of μ_t which leaves the components of $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ contained in the neighborhood of this equilibrium and returns back to these sets at a later time instant (as illustrated in

Figure 3-8). As in Step 4, Lipschitz continuity of ϕ is used to obtain an upper bound on the increase in the potential between the end points of this trajectory. Together with the lower bound obtained in Step 3, this provides a bound on how far μ_t can get from the component of $\mathcal{X}_{M\delta+\epsilon_2}$ contained in this neighborhood. Choosing ϵ_1 arbitrarily small (for a fixed ϵ_2) we obtain the tightest such bound. Using this result, we quantify how far μ_t can get from the equilibria of the game (after sufficient time) and the theorem follows.

Next we prove each of these steps.

Step 1: By definition $\mathcal{X}_{\alpha_1} \subset \mathcal{X}_\alpha$ for any $\alpha_1 \leq \alpha$. Since the feasible set of the maximization problem in (3.15) is given by \mathcal{X}_α , this implies that $f(\alpha_1) \leq f(\alpha)$, i.e., f is a weakly increasing function of its argument. Note that the feasible set of the maximization problem in (3.15) can be given by the correspondence $g(\alpha) = \mathcal{X}_\alpha$, which is upper semi continuous in α as shown in Lemma 3.5.2 (ii). Since as a function of \mathbf{x} , $\min_{k \in \{1, \dots, l\}} \|\mathbf{x} - \mathbf{x}_k\|$ is continuous it follows from Berge's maximum theorem (see Berge (1963)) that for $\alpha \geq 0$, $f(\alpha)$ is an upper semicontinuous function.

The set \mathcal{X}_0 corresponds to the set of equilibria of the game, hence $\mathcal{X}_0 = \{\mathbf{x}_1, \dots, \mathbf{x}_l\}$. Thus, the definition of f implies that $f(0) = 0$. Moreover, upper semicontinuity of f implies that for any $\epsilon > 0$, there exists some neighborhood V of 0, such that $f(x) \leq \epsilon$ for all $x \in V$. Since, $f(x) \geq 0$ by definition, this implies that $\lim_{x \rightarrow 0} f(x)$ exists and equals to 0.

Step 2: Let $\bar{\delta} > 0$ be small enough such that $M\bar{\delta} < a/2$. Since $\lim_{x \rightarrow 0} f(x) = 0$, it follows that for sufficiently small $\bar{\delta}$ and $\bar{\epsilon}$, we obtain $f(M\bar{\delta} + \bar{\epsilon}) < \frac{a\bar{\delta}}{48LM} < \frac{(a-M\bar{\delta})\bar{\delta}}{24LM}$ and $M\bar{\delta} + \bar{\epsilon} < a$.

Step 3: Let ϵ_1, ϵ_2 be such that $0 < \epsilon_1 < \epsilon_2$. Assume $T > 0$ is large enough so that for $t > T$,

$$\begin{aligned} \phi(\mu_{t+1}) - \phi(\mu_t) &\geq \frac{2\epsilon_1}{3(t+1)} \text{ if } \mu_t \notin \mathcal{X}_{M\delta+\epsilon_1}, \text{ and similarly} \\ \phi(\mu_{t+1}) - \phi(\mu_t) &\geq \frac{2\epsilon_2}{3(t+1)} \text{ if } \mu_t \notin \mathcal{X}_{M\delta+\epsilon_2}. \end{aligned} \quad (3.22)$$

Existence of T satisfying these inequalities follows from Lemma 3.5.3, since for large T and $t > T$, this lemma implies $\phi(\mu_{t+1}) - \phi(\mu_t) \geq \frac{\epsilon_1}{(t+1)} + O\left(\frac{1}{t^2}\right) \geq \frac{2\epsilon_1}{3(t+1)}$ if $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_1}$, and similarly if $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_2}$.

Since $\phi(\mu_t)$ increases outside $M\delta + \epsilon_1$ -equilibrium set for $t > T$, as (3.22) suggests, it follows that μ_t visits $\mathcal{X}_{M\delta+\epsilon_1}$ (and $\mathcal{X}_{M\delta+\epsilon_2}$ since $\mathcal{X}_{M\delta+\epsilon_1} \subset \mathcal{X}_{M\delta+\epsilon_2}$) infinitely often. Otherwise $\phi(\mu_t)$ increases unboundedly, and we reach a contradiction since mixed extension of the potential is a bounded function.

Assume that at some time after T , μ_t leaves $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ and returns back to $\mathcal{X}_{M\delta+\epsilon_1}$ at a later time instant. In this step, we quantify how much the potential increases when μ_t follows such a path. We first define time instants T_1, T_2, T'_1 , and T'_2 satisfying $T < T_1 \leq T_2 < T'_2 \leq T'_1$, as follows:

- T_1 is a time instant when μ_t leaves $\mathcal{X}_{M\delta+\epsilon_1}$, i.e., $\mu_{T_1-1} \in \mathcal{X}_{M\delta+\epsilon_1}$ and $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_1}$ for $T_1 \leq t < T'_1$.
- T_2 is a time instant when μ_t leaves $\mathcal{X}_{M\delta+\epsilon_2}$, i.e., $\mu_{T_2-1} \in \mathcal{X}_{M\delta+\epsilon_2}$ and $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_2}$ for $T_2 \leq t < T'_2$.
- T'_2 is the first time instant after T_2 when μ_t returns back to $\mathcal{X}_{M\delta+\epsilon_2}$, i.e., $\mu_{T'_2-1} \notin \mathcal{X}_{M\delta+\epsilon_2}$ and $\mu_{T'_2} \in \mathcal{X}_{M\delta+\epsilon_2}$.
- T'_1 is the first time instant after T_1 when μ_t returns back to $\mathcal{X}_{M\delta+\epsilon_1}$, i.e., $\mu_{T'_1-1} \notin \mathcal{X}_{M\delta+\epsilon_1}$ and $\mu_{T'_1} \in \mathcal{X}_{M\delta+\epsilon_1}$.

The definitions are illustrated in Figure 3-8. We next provide a lower bound on the quantity $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$. Note that if there are multiple time instants between T_1 and T'_1 for which μ_t leaves $\mathcal{X}_{M\delta+\epsilon_2}$ (as in the figure), any of these time instants can be chosen as T_2 to obtain a lower bound.




Figure 3-8: Trajectory of μ_t (initialized at the left end of the dashed line) is illustrated. T_1 and T_2 correspond to the time instants μ_t leaves $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ respectively. T'_1 and T'_2 correspond to the time instants μ_t enters $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ respectively.

By definition, for t such that $T_2 \leq t < T'_2$, we have $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_2}$, and for t such that $T_1 \leq t < T_2$ or $T'_2 \leq t < T'_1$, we have $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_1}$. Thus, it follows from (3.22) that

$$\phi(\mu_{t+1}) - \phi(\mu_t) \geq \frac{2\epsilon_2}{3(t+1)} \quad \text{for } T_2 \leq t < T'_2, \quad (3.23)$$

and consequently,

$$\phi(\mu_{T'_2}) - \phi(\mu_{T_2}) = \sum_{t=T_2}^{T'_2-1} \phi(\mu_{t+1}) - \phi(\mu_t) \geq \sum_{t=T_2}^{T'_2-1} \frac{2\epsilon_2}{3(t+1)}. \quad (3.24)$$

Similarly, since $\mu_t \notin \mathcal{X}_{M\delta+\epsilon_1}$ for t such that $T_1 \leq t < T_2$ or $T'_2 \leq t < T'_1$, using (3.22) we establish

$$\phi(\mu_{T'_1}) - \phi(\mu_{T'_2}) = \sum_{t=T'_2}^{T'_1-1} \phi(\mu_{t+1}) - \phi(\mu_t) \geq \sum_{t=T'_2}^{T'_1-1} \frac{2\epsilon_1}{3(t+1)}, \quad (3.25)$$

$$\phi(\mu_{T_2}) - \phi(\mu_{T_1}) = \sum_{t=T_1}^{T_2-1} \phi(\mu_{t+1}) - \phi(\mu_t) \geq \sum_{t=T_1}^{T_2-1} \frac{2\epsilon_1}{3(t+1)}. \quad (3.26)$$

Since $\phi(\mu_{T'_1}) - \phi(\mu_{T_1}) = \left(\phi(\mu_{T'_1}) - \phi(\mu_{T'_2})\right) + \left(\phi(\mu_{T'_2}) - \phi(\mu_{T_2})\right) + \left(\phi(\mu_{T_2}) - \phi(\mu_{T_1})\right)$, it follows from (3.24), (3.25) and (3.26) that

$$\phi(\mu_{T'_1}) - \phi(\mu_{T_1}) \geq \sum_{t=T_2}^{T'_2-1} \frac{2\epsilon_2}{3(t+1)}. \quad (3.27)$$

Step 4: Let $\epsilon_2 = a - M\bar{\delta}$, and $\epsilon_1 = \bar{\epsilon}$. By definition of $\bar{\epsilon}$ and $\bar{\delta}$ (see Step 2), it follows that $\epsilon_2 > \epsilon_1 > 0$. Assume that $\delta < \bar{\delta}$. Since $a = M\bar{\delta} + \epsilon_2 > M\delta + \epsilon_2 > M\delta + \epsilon_1$ we obtain $\mathcal{X}_{M\delta+\epsilon_1} \subset \mathcal{X}_{M\delta+\epsilon_2} \subset \mathcal{X}_a$. By definition of a , it follows that components of $\mathcal{X}_{M\delta+\epsilon_1}$ and $\mathcal{X}_{M\delta+\epsilon_2}$ are also contained in disjoint neighborhoods of equilibria. Hence, the definition of f suggests that if $\mathbf{x} \in \mathcal{X}_{M\delta+\epsilon_1}$ then $\|\mathbf{x}_k - \mathbf{x}\| \leq f(M\delta + \epsilon_1)$ (similarly if $\mathbf{x} \in \mathcal{X}_{M\delta+\epsilon_2}$, then $\|\mathbf{x}_k - \mathbf{x}\| \leq f(M\delta + \epsilon_2)$) for exactly one equilibrium \mathbf{x}_k .

Let T_1, T_2, T'_1 and T'_2 be defined as in Step 3. In this step, by obtaining an upper bound on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$ and refining the lower bound obtained in Step 3 for given values of ϵ_1 and ϵ_2 , we prove that after sufficient time μ_t can visit the component of $\mathcal{X}_{M\delta+\epsilon_1}$ in the neighborhood of a single equilibrium.

Assume that μ_t leaves the component of the $M\delta + \epsilon_1$ -equilibrium set in the neighborhood of equilibrium \mathbf{x}_k , and it reaches to another component in the neighborhood of equilibrium $\mathbf{x}_{k'}$. Since, by definition $\mu_{T_1-1}, \mu_{T'_1} \in \mathcal{X}_{M\delta+\epsilon_1}$, and $\mu_{T_2-1}, \mu_{T'_2} \in \mathcal{X}_{M\delta+\epsilon_2}$, it follows that μ_{T_1-1} and μ_{T_2-1} belong to neighborhoods of equilibrium \mathbf{x}_k , whereas, $\mu_{T'_1}$ and $\mu_{T'_2}$ belong to neighborhoods of $\mathbf{x}_{k'}$, i.e.,

$$\|\mathbf{x}_k - \mu_{T_1-1}\| \leq f(M\delta + \epsilon_1) \quad \text{and} \quad \|\mathbf{x}_k - \mu_{T_2-1}\| \leq f(M\delta + \epsilon_2), \text{ whereas,} \quad (3.28)$$

$$\|\mathbf{x}_{k'} - \mu_{T'_1}\| \leq f(M\delta + \epsilon_1) \quad \text{and} \quad \|\mathbf{x}_{k'} - \mu_{T'_2}\| \leq f(M\delta + \epsilon_2). \quad (3.29)$$

By definition of d we have $\|\mathbf{x}_k - \mathbf{x}_{k'}\| \geq d$. Since $a > M\delta + \epsilon_2$, it follows that $f(M\delta + \epsilon_2) < d/4$, and hence the second inequalities in (3.28) and (3.29) imply

$$\|\mu_{T'_2} - \mu_{T_2-1}\| > \frac{d}{2}. \quad (3.30)$$

Using this inequality, we next refine the lower bound on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$ obtained in Step 3. By (3.9), with an update at time t , the empirical frequency distribution can change

by at most

$$\|\mu_{t+1} - \mu_t\| = \frac{1}{t+1} \|\mu_t - I_t\| \leq \frac{1}{t+1} (\|\mu_t\| + \|I_t\|) \leq \frac{2M}{t+1}, \quad (3.31)$$

where the last inequality follows from the fact that $\mu_t = \{\mu_t^m\}_{m \in \mathcal{M}}$, and $I_t = \{I_t^m\}_{m \in \mathcal{M}}$, and $\|\mu_t^m\|, \|I_t^m\| \leq 1$, since $I_t^m, \mu_t^m \in \Delta E^m$. Hence, if T_2 is sufficiently large, then $\|\mu_{T_2} - \mu_{T_2-1}\|$ is small enough so that (3.30) implies $\|\mu_{T_2'} - \mu_{T_2}\| > \frac{d}{2}$. Using this together with (3.31), we conclude

$$\sum_{t=T_2}^{T_2'-1} \frac{2M}{t+1} \geq \sum_{t=T_2}^{T_2'-1} \|\mu_{t+1} - \mu_t\| \geq \left\| \sum_{t=T_2}^{T_2'-1} \mu_{t+1} - \mu_t \right\| = \|\mu_{T_2'} - \mu_{T_2}\| > \frac{d}{2}. \quad (3.32)$$

Thus, the lower bound on $\phi(\mu_{T_1'}) - \phi(\mu_{T_1})$ provided in (3.27) takes the following form:

$$\phi(\mu_{T_1'}) - \phi(\mu_{T_1}) \geq \sum_{t=T_2}^{T_2'-1} \frac{2\epsilon_2}{3(t+1)} \geq \frac{\epsilon_2 d}{6M}. \quad (3.33)$$

Next we provide an upper bound on $\phi(\mu_{T_1'}) - \phi(\mu_{T_1})$, using Lipschitz continuity of the potential and the properties of the $M\delta + \epsilon_1$ equilibrium set. Let $\bar{\phi}_k = \max_{\{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_k\| \leq f(M\delta + \epsilon_1)\}} \phi(\mathbf{x})$, and define \mathbf{y}_k as a strategy profile which achieves this maximum. Similarly, let $\underline{\phi}_{k'} = \min_{\{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_{k'}\| \leq f(M\delta + \epsilon_1)\}} \phi(\mathbf{x})$ and define $\mathbf{y}_{k'}$ as a strategy profile which achieves this minimum. Observe that

$$\begin{aligned} \underline{\phi}_{k'} - \bar{\phi}_k &= \phi(\mathbf{y}_{k'}) - \phi(\mathbf{y}_k) \\ &= \left(\phi(\mathbf{y}_{k'}) - \phi(\mu_{T_1'}) \right) + \left(\phi(\mu_{T_1'}) - \phi(\mu_{T_1}) \right) + \left(\phi(\mu_{T_1}) - \phi(\mathbf{y}_k) \right). \end{aligned} \quad (3.34)$$

Note that by (3.28) and (3.29), and the definitions of \mathbf{y}_k and $\mathbf{y}_{k'}$, we have $\mu_{T_1'}, \mathbf{y}_{k'} \in \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_{k'}\| \leq f(M\delta + \epsilon_1)\}$, and $\mu_{T_1-1}, \mathbf{y}_k \in \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_k\| \leq f(M\delta + \epsilon_1)\}$. Hence, using Lipschitz continuity of ϕ (and denoting the Lipschitz constant by L) it follows that $\phi(\mathbf{y}_{k'}) - \phi(\mu_{T_1'}) \geq -2Lf(M\delta + \epsilon_1)$, and $\phi(\mu_{T_1-1}) - \phi(\mathbf{y}_k) \geq -2Lf(M\delta + \epsilon_1)$. Moreover, (3.31) and Lipschitz continuity of ϕ imply that $\phi(\mu_{T_1}) - \phi(\mu_{T_1-1}) = O\left(\frac{1}{T_1}\right)$. Thus, using (3.34) we obtain the following upper bound on $\phi(\mu_{T_1'}) - \phi(\mu_{T_1})$:

$$\underline{\phi}_{k'} - \bar{\phi}_k + 4Lf(M\delta + \epsilon_1) + O\left(\frac{1}{T_1}\right) \geq \phi(\mu_{T_1'}) - \phi(\mu_{T_1}). \quad (3.35)$$

Using the lower and upper bounds we obtained in (3.33) and (3.35), it follows that

$$\underline{\phi}_{k'} - \bar{\phi}_k + 4Lf(M\delta + \epsilon_1) + O\left(\frac{1}{T_1}\right) \geq \frac{\epsilon_2 d}{6M}. \quad (3.36)$$

Since $\epsilon_2 = a - M\bar{\delta}$, and $\epsilon_1 = \bar{\epsilon}$, using the fact that f is an increasing function and $\delta < \bar{\delta}$,

it follows from (3.36) that

$$\underline{\phi}_{k'} - \bar{\phi}_k \geq \frac{(a - M\bar{\delta})d}{6M} - 4Lf(M\bar{\delta} + \bar{\epsilon}) + O\left(\frac{1}{T_1}\right) \geq \frac{(a - M\bar{\delta})d}{6M} - 4Lf(M\bar{\delta} + \bar{\epsilon}) + O\left(\frac{1}{T_1}\right)$$

Note that (3.21) implies $\frac{(a - M\bar{\delta})d}{6M} - 4Lf(M\bar{\delta} + \bar{\epsilon}) > 0$. Thus, for sufficiently large T_1 we obtain $\underline{\phi}_{k'} - \bar{\phi}_k > 0$. Therefore, we conclude when μ_t leaves the component of $\mathcal{X}_{M\delta + \epsilon_1}$ contained in the neighborhood of some equilibrium \mathbf{x}_k , and enters that of another equilibrium $\mathbf{x}_{k'}$, then the minimum potential in the new neighborhood is strictly larger than the maximum potential in the older one (for sufficiently large T_1). Since this is true for arbitrary equilibria \mathbf{x}_k and $\mathbf{x}_{k'}$, it follows that after entering the component of $\mathcal{X}_{M\delta + \epsilon_1}$ in the neighborhood of $\mathbf{x}_{k'}$, μ_t cannot return to the component in the neighborhood of \mathbf{x}_k , as doing so contradicts with the relation between the minimum and maximum potentials in these neighborhoods. Thus, after sufficient time, μ_t can visit the component of $\mathcal{X}_{M\delta + \epsilon_1}$ (or equivalently $\mathcal{X}_{M\delta + \bar{\epsilon}}$) in the neighborhood of a single equilibrium.

Step 5: Let ϵ_1 , and ϵ_2 be such that $0 < \epsilon_1 < \epsilon_2 \leq \bar{\epsilon}$. As established in Step 4, there exists some T , such that for $t > T$, μ_t visits the component of $\mathcal{X}_{M\delta + \bar{\epsilon}}$, in the neighborhood of a single equilibrium, say \mathbf{x}_k .

Assume that T_1, T_2, T'_1 and T'_2 are defined as in Step 3, and let $T_1 > T + 1$. Since $\epsilon_1 < \epsilon_2 \leq \bar{\epsilon}$, we have $\mathcal{X}_{M\delta + \epsilon_1} \subset \mathcal{X}_{M\delta + \epsilon_2} \subset \mathcal{X}_{M\delta + \bar{\epsilon}}$, and $T_1 > T + 1$ implies that μ_t can only visit the components of $\mathcal{X}_{M\delta + \epsilon_1}$ and $\mathcal{X}_{M\delta + \epsilon_2}$ contained in the neighborhood of \mathbf{x}_k . Following a similar approach to Step 4, we next obtain upper and lower bounds on $\phi(\mu_{T'_1}) - \phi(\mu_{T_1})$, and use these bounds to establish convergence to the mixed equilibrium set given in the theorem statement.

Define d^* as the maximum distance of μ_t from $\mathcal{X}_{M\delta + \epsilon_2}$ for t such that $T + 1 < T_2 \leq t \leq T'_2 - 1$, i.e.,

$$d^* = \max_{\{t | T_2 \leq t \leq T'_2 - 1\}} \min_{\mathbf{x} \in \mathcal{X}_{M\delta + \epsilon_2}} \|\mu_t - \mathbf{x}\|.$$

Since $\mu_{T_2 - 1}, \mu_{T'_2} \in \mathcal{X}_{M\delta + \epsilon_2}$ by definition, the total length of the trajectory between $T_2 - 1$ and T'_2 is an upper bound on $2d^*$, i.e.,

$$2d^* \leq \sum_{t=T_2-1}^{T'_2-1} \|\mu_{t+1} - \mu_t\|.$$

As explained in (3.31), $\|\mu_{t+1} - \mu_t\| \leq \frac{2M}{t+1}$, thus the above inequality implies

$$2d^* \leq \sum_{t=T_2-1}^{T'_2-1} \frac{2M}{t+1} = \sum_{t=T_2}^{T'_2-1} \frac{2M}{t+1} + \frac{2M}{T_2}. \quad (3.37)$$

Using this inequality, the lower bound in (3.27) implies

$$\phi(\mu_{T_1'}) - \phi(\mu_{T_1}) \geq \sum_{t=T_2}^{T_2'-1} \frac{2\epsilon_2}{3(t+1)} \geq \left(d^* - \frac{M}{T_2}\right) \frac{2\epsilon_2}{3M} \quad (3.38)$$

We next obtain an upper bound on $\phi(\mu_{T_1'}) - \phi(\mu_{T_1})$. By definition of f , $\mathcal{X}_{M\delta+\epsilon_1}$ is contained in $f(M\delta + \epsilon_1)$ neighborhoods of equilibria. For $T_1 > T + 1$, μ_t can only visit the component of $\mathcal{X}_{M\delta+\epsilon_1}$ in the neighborhood of \mathbf{x}_k , as can be seen from the definition of T . Thus, since $\mu_{T_1-1}, \mu_{T_1'} \in \mathcal{X}_{M\delta+\epsilon_1}$, it follows that $\mu_{T_1-1}, \mu_{T_1'} \in \{\mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_k\| \leq f(M\delta + \epsilon_1)\}$. By Lipschitz continuity of the potential function it follows that $\phi(\mu_{T_1'}) - \phi(\mu_{T_1-1}) \leq 2f(M\delta + \epsilon_1)L$. Additionally, by (3.31) Lipschitz continuity also implies that $\phi(\mu_{T_1}) - \phi(\mu_{T_1-1}) \leq \frac{2ML}{T_1}$. Combining these we obtain the following upper bound on $\phi(\mu_{T_1'}) - \phi(\mu_{T_1})$:

$$\phi(\mu_{T_1'}) - \phi(\mu_{T_1}) \leq 2f(M\delta + \epsilon_1)L + \frac{2ML}{T_1}. \quad (3.39)$$

It follows from the upper and lower bounds on $\phi(\mu_{T_1'}) - \phi(\mu_{T_1})$ given in (3.38) and (3.39) that

$$\left(d^* - \frac{M}{T_2}\right) \frac{2\epsilon_2}{3M} \leq 2f(M\delta + \epsilon_1)L + \frac{2ML}{T_1}$$

Thus, for sufficiently large T_1 (and hence T_2), we obtain

$$d^* \leq \frac{3f(M\delta + \epsilon_1)ML}{\epsilon_2} + \frac{3M^2L}{\epsilon_2 T_1} + \frac{M}{T_2} \leq \frac{4f(M\delta + \epsilon_1)ML}{\epsilon_2}. \quad (3.40)$$

Note that in the above derivation ϵ_1 is an arbitrary number that satisfies $0 < \epsilon_1 < \epsilon_2$. Thus, (3.40) implies that

$$d^* \leq \limsup_{\epsilon_1 \rightarrow 0} \frac{4f(M\delta + \epsilon_1)ML}{\epsilon_2} \leq \frac{4f(M\delta)ML}{\epsilon_2}, \quad (3.41)$$

where the last inequality follows by upper semicontinuity of f . Thus, by definition of d^* , we conclude that μ_t converges d^* neighborhood of $\mathcal{X}_{M\delta+\epsilon_2}$. Hence, using (3.41), we can establish convergence of μ_t to

$$\left\{ \mathbf{x} \mid \|\mathbf{x} - \mathbf{y}\| \leq \frac{4f(M\delta)ML}{\epsilon_2}, \text{ for some } \mathbf{y} \in \mathcal{X}_{M\delta+\epsilon_2} \right\}. \quad (3.42)$$

Observe that definition of f implies if $\mathbf{y} \in \mathcal{X}_{M\delta+\epsilon_2}$, then for some equilibrium \mathbf{x}_k we have $\|\mathbf{x}_k - \mathbf{y}\| \leq f(M\delta + \epsilon_2)$. Thus, using (3.42) and triangle inequality, we conclude that μ_t converges to

$$\left\{ \mathbf{x} \mid \|\mathbf{x} - \mathbf{x}_k\| \leq \frac{4f(M\delta)ML}{\epsilon_2} + f(M\delta + \epsilon_2), \text{ for some equilibrium } \mathbf{x}_k \right\}. \quad (3.43)$$

Noting that in (3.43) ϵ_2 is an arbitrary number satisfying $0 < \epsilon_2 \leq \bar{\epsilon}$, the theorem follows. \square

Part II

Iterative Auction Design for Graphical Valuations

Chapter 4

Tree Valuations

4.1 Introduction and Background

Iterative auctions are a class of mechanisms that are commonly employed in practice. In these auctions, the auctioneer sets prices for the items she is selling, bidders report which items they are interested in at the given prices, and in response to these reports the auctioneer updates the prices. This process terminates when the auctioneer determines a final allocation of items to the bidders. Examples of iterative auctions include the auctions used for selling art, antiquities, wine, jewelry, electricity, natural resources, bus routes, spectrum, and the auctions used for procurement.¹

The well-known English and Dutch auctions can be viewed as examples of single-item iterative auctions. When bidders have independent private values, these auctions allocate the item efficiently, i.e., the bidder with the highest value receives the item. On the other hand, in more general multi-item settings (such as spectrum or procurement auctions) the iterative auctions that are present in the literature do not always have similar efficiency guarantees. More precisely, they either implement the efficient outcome under restrictive assumptions (such as the gross substitutes assumption (Gul and Stacchetti, 2000; Ausubel, 2006)), or they require complicated pricing structures that involve a different price for each bundle of items (Ausubel and Milgrom, 2002; Bikhchandani et al., 2002; Vohra, 2011). The auction formats in the first category do not allow for value complementarity between different items, which is commonly observed in various practical auction environments. Those in the second category may not be practical since they require the number of different prices that are reported to the bidders at each stage of the auction to be exponential in the number of items.

Motivated by these considerations, in this part of the thesis, we study the question of

¹For specific practical examples of these auctions, see the websites of auction houses such as Sotheby's (<http://www.sothebys.com/>) and Christie's (<http://www.christies.com/>), online retailers such as eBay (<http://www.ebay.com/>), and auction service providers such as Power Auctions LLC (<http://www.powerauctions.com/>).

iterative auction design for multi-item environments. Our main contribution is to develop simple iterative auction formats for settings that involve both value complementarity and substitutability. In the auction formats we provide, bidders dynamically update their bids in the course of the auction. Following a similar approach to that of the first part of this thesis, we establish that the strategy updates of bidders converge to the efficient outcome. Moreover, the auction formats we provide rely on a simple pricing rule.

We obtain our results by focusing on a special class of valuation functions, which we refer to as graphical valuations. Graphical valuations allow for a compact representation of the value functions of the bidders. In particular, value functions that belong to this class are associated with a value graph, nodes of which correspond to the items that are being sold by the auctioneer. There are edges between items that can exhibit value complementarity or substitutability. We associate weights with the nodes and edges of the underlying graph. Positive weights associated with an edge capture value complementarity between the nodes (items) at the end points of this edge, and negative weights capture substitutability. The value a player has for a set of items is equal to the sum of the node weights and the edge weights of the subgraph obtained by restricting the original value graph to this set of nodes (items).

It can be seen that graphical valuations are not fully general, i.e., there are value functions that cannot be represented as a graphical valuation. On the other hand, we believe that this class captures the value complementarity/substitutability in many practical auction settings reasonably well. For instance, consider the spectrum auctions, where the items that are auctioned correspond to spectrum bands at different regions, and there is value complementarity/substitutability between adjacent regions. In this setting, the values of bidders can naturally be captured with graphical valuations, using a graph that has a node for each region. There are edges in this graph between nodes representing adjacent regions, and the associated weights capture the value complementarity/substitutability between these regions.

The standard approach in the literature (see Vohra (2011)) for developing iterative auctions involves three main steps. Efficient iterative auctions implicitly solve an optimization problem, and find the welfare maximizing allocation. Before we develop iterative auctions for graphical valuations, we first focus on this optimization problem, and formulate it as a linear program. Second, we consider iterative algorithms that can be used for solving linear programs, and apply these to the solution of the linear programming formulation of the efficient allocation problem. Finally, we establish that these algorithms suggest natural iterative auction formats that converge to the efficient outcome, when bidders truthfully reveal their demand. By charging bidders appropriate final payments, we show that it is an equilibrium for bidders to truthfully reveal their demand in these auctions.

The existing iterative auction formats that follow this approach and allow for complementarity in valuations, rely on exponentially many prices for implementing the efficient

outcome (Bikhchandani et al., 2002; De Vries et al., 2007; Mishra and Parkes, 2007; Vohra, 2011). In contrast, our main contribution in this part of the thesis is to develop efficient iterative auction formats that rely on simple pricing rules. In order to accomplish this, we follow the above outlined approach, and by exploiting the structure of graphical valuations, we first obtain simpler linear optimization formulations for the efficient allocation problem than the ones present in the literature. Then, using iterative solutions of these new formulations, we develop novel iterative auction formats that rely on simple pricing rules, and guarantee efficiency for graphical valuations.

In this chapter and Chapter 5 we focus on the first component of the iterative auction design framework, and provide linear programming formulations that can be solved to identify the efficient outcome. We defer the question of iterative auction design using these LP formulations to Chapter 6. Additionally, in Chapter 6, we also address how to design payment schemes that guarantee that the efficient outcome is implemented at an (ex-post perfect) equilibrium of our auctions.

We start our exposition in this chapter by formally introducing our valuation model, and discussing some structural properties of iterative auctions. An important component of iterative auction design is the choice of the pricing rule used for running auctions. In this chapter, we present an important pricing rule, anonymous item pricing, that is commonly used in the literature for the design of iterative auctions (Ausubel, 2006). We also discuss a natural termination condition for iterative auctions that rely on this pricing rule: the auctioneer terminates the auction when a “market clearance” condition holds, i.e., when all bidders demand disjoint sets of items, and all items are demanded by some bidder. At such an outcome, no bidder needs not compete with the remaining bidders to acquire the set of items that she demands (since the demand sets are disjoint), thereby making this outcome a natural termination point for the auction. Moreover, this termination condition is equivalent to convergence of the iterative auction to a Walrasian equilibrium. Hence, it is possible to design iterative auction formats that rely on anonymous item pricing and the aforementioned termination condition if and only if a Walrasian equilibrium exists.

The main contribution of this chapter is to establish that when the underlying value graph is a tree, and satisfies an additional technical (sign consistency) condition, a Walrasian equilibrium exists. Intuitively, the sign consistency condition suggests that if a bidder views two items as complements (substitutes), so do the remaining bidders. It is known that the existence of a Walrasian equilibrium is equivalent to existence of integral optimal solutions to a linear programming formulation of the efficient allocation problem. Thus, our result immediately leads to a linear program that can be solved to identify the efficient allocation for sign-consistent tree valuations. Moreover, as we demonstrate in Chapter 6, iterative solutions of this LP formulation can be used to obtain iterative auction formats that terminate when a Walrasian equilibrium is identified, and allocates items to bidders efficiently, when the underlying valuations are sign-consistent tree valuations.

We also demonstrate that if we relax the sign consistency assumption, or the tree assumption, solving this linear programming formulation no longer gives the efficient outcome and a Walrasian equilibrium does not exist. This suggests that for more general graphical valuations, it is not possible to implement the efficient outcome using iterative auction formats that rely on anonymous item pricing, and terminate when a market clearance condition holds. Interestingly, by considering more general pricing rules and linear programming formulations, the efficient allocation can be implemented using iterative auction formats that rely on a similar termination condition. We defer the discussion of such general pricing rules and linear programming formulations to Chapter 5.

4.1.1 Related literature

In standard multi-item auction settings, the VCG mechanism can be used to implement the efficient outcome in dominant strategy equilibria. Despite this desirable strategic feature, VCG mechanisms are rarely used in practice. On the other hand, iterative auction formats which share similar equilibrium properties to VCG mechanisms are prevalent (see for instance Ausubel and Milgrom (2006); Rothkopf et al. (1990); Engelbrecht-Wiggans and Kahn (1991)), and have found applications in spectrum auctions, electricity auctions, online markets (such as eBay) (McAfee et al., 2010; Ausubel and Cramton, 2004; Ausubel, 2004), as well as procurement settings (Hohner et al., 2003; Cramton et al., 2006). This has stimulated significant interest in recent literature, and led to development of a number of novel multi-item iterative auction formats. Examples include the package bidding auction (Ausubel and Milgrom, 2002), clinching auction and its variants (Ausubel, 2004, 2006), auctions that rely on universally competitive equilibria (UCE) (Mishra and Parkes, 2007), and best response auction of Nisan et al. (2011a).

Many of the iterative auctions present in the literature implement the efficient outcome under the restrictive gross substitutes assumptions. Intuitively, gross substitute valuations suggest that if a bidder demands a set of items at given prices, and the price of one of these items increases, the demand for the remaining items cannot decrease (Kelso and Crawford, 1982; Gul and Stacchetti, 1999). Under the gross substitutes property, a Walrasian equilibrium exists and the prices that support it have a lattice structure. These results can be exploited to define simple tâtonnement processes (and auction formats) that converge to a Walrasian equilibrium and an efficient allocation (Gul and Stacchetti, 2000; Ausubel, 2006). The gross substitutes property, on the other hand, does not allow for any value complementarity between different items, which is a key feature of important combinatorial auction settings. A generalization of this class, which allows for a very specific value complementarity structure, is the class of gross substitutes and complements (GSC), see Sun and Yang (2006, 2009). The GSC structure suggests that items can be grouped into two sets so that all items in a given set are gross substitutes, and items that belong to different sets are complements. It is possible to establish the existence of a Walrasian equilibrium

and provide simple iterative auctions (Sun and Yang, 2006, 2009) for such valuations. However, these results are limited to the particular complementarity structure imposed by the GSC valuations, and do not overlap with our contributions (see Section 4.6 for a detailed discussion).

The problem of finding the efficient outcome in a combinatorial setting is hard both from a computational complexity and a communication complexity point of view (Lehmann et al., 2006; Nisan and Segal, 2006; Cramton et al., 2006; Blumrosen and Nisan, 2010). This motivated considering classes of valuation functions with additional structure (Blumrosen and Nisan, 2010; Cramton et al., 2006). Recently, Zhou et al. (2009) and Abraham et al. (2012) considered graphical valuation structures that are similar to those we consider in this thesis. In these papers, authors characterize the complexity of auction design for (hyper) graphical valuations, and develop approximately efficient auctions for settings that do not exhibit substitutabilities. In contrast, in this part of the thesis, we adopt a similar value model to the one present in these papers, but develop efficient iterative auctions for valuations that allow for both complementarities and substitutabilities.

Efficient iterative auctions implicitly solve an optimization problem: they find a welfare-maximizing allocation of goods to bidders. In many settings, it is possible to formulate the underlying efficient allocation problem as a linear program (with possibly exponentially many variables) (Bikhchandani and Mamer, 1997; Bikhchandani and Ostroy, 2002). Bikhchandani and Mamer (1997) establish that a particular linear programming formulation, which we also employ in this work, has optimal solutions that are integral, and can be used to find the efficient allocation, if and only if a Walrasian equilibrium exists. However, the associated integrality results present in the literature are restricted to settings, where the gross substitutes condition holds. Interestingly, our results suggest that for sign-consistent tree valuations, this LP formulation has an optimal solution that is integral, this solution corresponds to an efficient allocation, and a Walrasian equilibrium exists.

The idea of developing iterative auctions by (i) first formulating the efficient allocation problem as a linear program, (ii) then obtaining iterative algorithms for solutions of these problems, and interpreting them as iterative auction formats, and (iii) guaranteeing incentive compatibility by complementing these algorithms with appropriate payment schemes is present in the earlier literature (Parkes, 1999; Bikhchandani et al., 2002; De Vries and Vohra, 2003; Parkes, 2006; De Vries et al., 2007; Vohra, 2011). In the auctions present in these papers, the auctioneer searches for the efficient outcome by modifying the price or a temporary allocation of goods to bidders over time. In response to the price/allocation change, bidders update their demand. When bidders truthfully reveal their demand, such procedures correspond to applications of certain iterative algorithms to the solution of the linear programming formulation of the efficient allocation problem. In order to guarantee incentive compatibility and truthful demand revelation, the auctioneer uses the information revealed in the course of the auction and computes final payments for bidders (that are of-

ten equal to the VCG payments). Some of the auction formats that rely on this approach implement the efficient allocation in settings that are more general than the gross substitutes case, e.g., (Bikhchandani et al., 2002; De Vries et al., 2007; Ausubel and Milgrom, 2002; Mishra and Parkes, 2007; Vohra, 2011). However, the iterative algorithms (and the corresponding iterative auctions) present in these papers rely on using exponentially many prices at each step, as the underlying LP formulations do not admit a simple structure. This leads to a serious drawback for implementation, since at each step the number of prices the auctioneer needs to communicate to the bidders is exponential in the number of the items. Consequently, this pricing rule is prohibitive in multi-item environments where more than a few items are auctioned together. In contrast, our objective in this work is to develop efficient iterative auction formats that rely on simple pricing rules. We accomplish this by following a similar approach to the one outlined above in the context of graphical valuations. Importantly, by exploiting the special structure of graphical valuations, we obtain simpler linear programming formulations of the efficient allocation problem than the ones present in the literature (that lead to offering exponentially many prices to bidders), and novel efficient iterative auction formats that rely on simple pricing rules.

Finally, a related focus to ours is present in Bikhchandani et al. (2011), where an iterative auction format that employs only a single price, and guarantees efficiency in settings with an additional combinatorial structure (a matroid structure) is provided. We note that these results are not directly applicable in our setting, since the valuations we consider do not exhibit the required structure.

4.1.2 Outline

The rest of this chapter is organized as follows: In Section 4.2, we introduce the model and notation that will be used in this part of the thesis. Additionally, we discuss the Walrasian equilibrium concept and provide a condition for its existence. In Section 4.3, we focus on cases where the underlying value graph is a (sign-consistent) tree graph. We show that a Walrasian equilibrium exists for this class of valuation functions, and this leads to a linear programming formulation for obtaining the efficient allocation. In Section 4.4, we show that the solution of this LP formulation does not give the efficient outcome, if we relax the sign consistency or tree valuation assumptions. This suggests that for general graphical valuations more complex LP formulations may be necessary for obtaining the efficient outcome. We defer the discussion of these more general LP formulations and the iterative auctions that rely on iterative solutions of these LPs to Chapters 5 and 6. A brief summary of the main contributions of this chapter is provided in Section 4.5. We provide a comparison of our results with a well-studied class of valuation functions that have a Walrasian equilibrium, and exhibit both value complementarity and substitutability in Section 4.6. Some of the technical proofs are delegated to Section 4.7.

4.2 Model and Preliminaries

In this section we first describe the valuation model that we focus on in this part of the thesis (Section 4.2.1). Then, in Section 4.2.2 we discuss pricing rules and termination conditions that can be used for the design of iterative auctions. In particular, we focus on a natural termination condition that closely relates to the Walrasian equilibrium concept. Additionally, we present necessary and sufficient conditions for the existence of a Walrasian equilibrium.

4.2.1 Graphical Valuations

In this part of the thesis, we focus on settings where an auctioneer sells a number of (heterogeneous) items to a finite set of bidders. In particular, our setting involves

- A finite set of items \mathcal{N} , such that $|\mathcal{N}| = N$,
- A finite set of bidders \mathcal{M} , such that $|\mathcal{M}| = M$,
- Value functions $v^m : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$ for each player m , such that for a given set of items $S \subset \mathcal{N}$, $v^m(S) \geq 0$ captures the value of player $m \in \mathcal{M}$.

We make two standard assumptions about the valuation functions: bidders have zero value for receiving no items, and bidders have weakly larger valuations for larger bundles.

Assumption 4.2.1. *We assume that $v^m(\emptyset) = 0$, i.e., bidders have value zero, for not receiving any items. Additionally, we assume that bidders have monotone increasing valuations, i.e., $v^m(S_1) \leq v^m(S_2)$ if $S_1 \subset S_2$.*

The value a player has for a set S need not be additive. That is, it may be the case that $v^m(S) \neq \sum_{i \in S} v^m(\{i\})$. If items i and j are such that $v^m(\{ij\}) \geq v^m(\{i\}) + v^m(\{j\})$, then we say that these items are *pairwise complementary*. On the other hand, if i and j are such that $v^m(\{ij\}) \leq v^m(\{i\}) + v^m(\{j\})$, we refer to them as *pairwise substitutes*. In this work we are mainly interested in pairwise complementarity/substitutability, and unless noted otherwise, we refer to pairwise complementarity/substitutability simply as complementarity/substitutability. We provide a discussion of the differences between graphical valuations and other related special valuation classes such as gross substitutes, gross substitutes and complements (Kelso and Crawford, 1982; Gul and Stacchetti, 1999; Sun and Yang, 2006); sub/superadditive valuations, and sub/supermodular valuations (Blumrosen and Nisan, 2010) in Section 4.6.

In this thesis, we impose additional structure on the valuation functions. In particular, we assume that the value functions admit a compact graphical representation. Before, we explain this additional structure, we introduce the notion of *value graph*.

Definition 4.2.1 (Value Graph). Let $G = (\mathcal{N}, E)$ be a graph such that the set of nodes corresponds to the set of items \mathcal{N} and there are edges between nodes (items) that may exhibit value complementarity and substitutability. We refer to G as a value graph for set of items \mathcal{N} .

We next use the notion of value graph to introduce the notion of graphical valuations.

Definition 4.2.2 (Graphical Valuations). Let $G = (\mathcal{N}, E)$ be a given value graph. We say that the value function $v : 2^{\mathcal{N}} \rightarrow \mathbb{R}^+$ is a graphical valuation (with respect to G) if:

- there exist positive node weights $w_i \geq 0$ for each $i \in \mathcal{N}$
- there exist (positive or negative) edge weights w_{ij} for each $(i, j) \in E$
- v is such that $v(S) = \sum_{i \in S} w_i + \sum_{(i,j) \in E | i,j \in S} w_{ij}$.

We refer to the weights $\{w_{ij}\} \cup \{w_i\}$ as weights consistent with G .

This definition implies that a valuation function is graphical, if there exists node and edge weights associated with the underlying value graph, such that the value of any bundle S equals to the sum of the node weights and edge weights, for nodes and edges contained in an induced subgraph of G with set of nodes S . For an example see Figure 4-1.

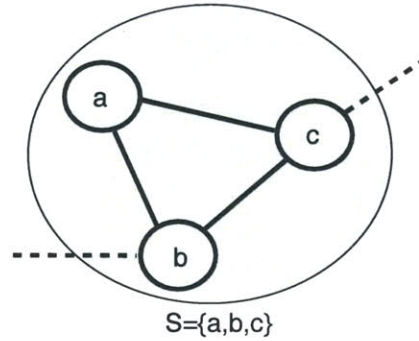


Figure 4-1: For a graphical valuation v , the value of bundle $S = \{a, b, c\}$ can be given as $v(S) = w_a + w_b + w_c + w_{ab} + w_{ac} + w_{bc}$.

Not all valuation functions are graphical valuations. In full generality, a value function associates a value with each bundle of items, and hence can be thought of as a vector of length $2^{\mathcal{N}}$. On the other hand, the definition of graphical valuations suggest that these valuations can be uniquely defined by specifying N node weights and at most N^2 edge weights. This implies that the set of graphical valuations has smaller dimension than the set of general valuation functions, and hence is not fully general. On the other hand, graphical valuations can be compactly represented by specifying the edge/node weights for the corresponding value graph.

Despite not being fully general, the graphical valuations can naturally capture the pairwise complementarity and substitutability in valuations. For instance, assume that i and j are two items such that for graphical valuation v we have $w_{ij} \geq 0$. Then it can be seen that $v(\{i\}) + v(\{j\}) = w_i + w_j \leq w_i + w_j + w_{ij} = v(\{i, j\})$, and hence i, j are pairwise complementary items. Conversely, if $w_{ij} \leq 0$, then $v(\{i\}) + v(\{j\}) = w_i + w_j \geq w_i + w_j + w_{ij} = v(\{i, j\})$ and i, j are pairwise substitutes.

In this work, we assume that all players have graphical valuations:

Assumption 4.2.2. *There exists a value graph $G = (\mathcal{N}, E)$ such that the valuation function of each player is a graphical valuation with respect to G . That is, for each player $m \in \mathcal{M}$, there exists weights*

- w_i^m for each node $i \in \mathcal{N}$, and
- w_{ij}^m for each edge $(i, j) \in E$,

such that $v^m(S) = \sum_{i \in S} w_i^m + \sum_{(i,j) \in E | i,j \in S} w_{ij}^m$.

Observe that in the above definition, the assumption that the valuation functions of all players are with respect to the same value graph is without loss of generality. To see this, assume that the valuation function of each player $m \in \mathcal{M}$ is consistent with some value graph $G^m = (\mathcal{N}, E^m)$. It follows from Definition 4.2.2 that the value function of each player is also consistent with the graph $G = (\mathcal{N}, \cup_m E^m)$, and a set of weights \hat{w}^m such that

$$\hat{w}_{ij}^m = \begin{cases} w_{ij}^m & \text{if } (i, j) \in E^m \\ 0 & \text{otherwise.} \end{cases}$$

Thus, even if the valuation function of each bidder is derived from a different value graph, it is possible to find a value graph and (bidder-specific) weights such that the valuation functions are also consistent with this value graph.

We believe that graphical valuations appropriately capture the value complementarity/substitutability in many practical auction settings, such as spectrum auctions, truck route auctions, and real estate auctions. In these settings, the items that are auctioned correspond to different geographical regions, and complementarities and substitutabilities are between neighboring geographical regions. For instance, for spectrum auctions, complementarities between adjacent geographical regions are present (Cramton et al., 1997; Moreton and Spiller, 1998), due to considerations such as roaming and interference. Similarly, different bands in the same geographical region can be viewed as substitutes, as the bidders may only have limited demand for spectrum in each geographical region. Such complementarities and substitutabilities can naturally be captured by graphical valuations by associating a node with each spectrum band - geographical region pair, and an edge with pairs of spectrum bands in adjacent (or the same) geographical regions (see Figure 4-2).




Figure 4-2: Consider a spectrum auction where two bands (A & B), over six geographical regions are sold. Agents view the bands in neighboring geographical regions as complements, while they view different bands in the same geographical region as substitutes. This can be captured using the graphical model in the above figure, and assigning positive weights to the solid lines, and negative weights to the dashed ones.

When the true value structure cannot be captured by graphical valuations, it is possible to consider their generalizations. In particular, assume that we associate a weight with each k -clique of the underlying value graph for $k \in \{1, \dots, r\}$, and the value a player has for a bundle S is given by the sum of the weights of all k -cliques ($k \in \{1, \dots, r\}$), nodes of which are contained in S . It can be seen that the graphical valuations defined above correspond to the case of $r = 2$. Choosing larger values of r allows for generalizing graphical valuations, and capturing more complicated structures of value complementarity and substitutability.² In particular, it can be shown that it is possible to represent any valuation function using generalized graphical valuations, by choosing the underlying value graph as a complete graph, and associating a weight with all cliques of the underlying graph (i.e., the $r = N$ case). For simplicity, in this chapter we focus on graphical valuations (i.e., the $r = 2$ case). In the next chapter, we will also discuss how our approach extends to the more general setting where $r > 2$ (see Section 5.5).

We close this section by formally defining the objective of our iterative auctions. We first introduce some relevant definitions.

Definition 4.2.3 (Feasible allocation). *Given sets $S^m \subset \mathcal{N}$, for all $m \in \mathcal{M}$, we say that $\{S^m\}_{m \in \mathcal{M}}$ is a feasible allocation if*

- *Each player $m \in \mathcal{M}$ receives a set of items $S^m \subset \mathcal{N}$,*
- *Each item is assigned to at most one player, and hence $S^m \cap S^l = \emptyset$ for $m \neq l$, and $m, l \in \mathcal{M}$.*

²In (Abraham et al., 2012), motivated by similar considerations, authors focus on hyper-graph valuations, characterize the computational complexity of auction design, and provide approximately efficient sealed bid auctions for hyper-graph valuations that exhibit only complementarities.

We say that a feasible allocation is complete if for every item i , there exists a bidder m such that $i \in S^m$. We denote the set of complete feasible allocations by χ .

An efficient allocation is a feasible allocation $\{S^m\}_{m \in \mathcal{M}}$ that maximizes the total value, i.e.,

$$\begin{aligned} \sum_m v^m(S^m) &= \max_{\{Z^m\}_m} \sum_m v^m(Z^m) \\ \text{s.t.} \quad & Z^m \subset \mathcal{N}, \text{ for all } m \in \mathcal{M} \\ & Z^m \cap Z^l = \emptyset, \text{ for all } m, l \in \mathcal{M}, m \neq l. \end{aligned}$$

Note that under Assumption 4.2.1 there exists an efficient allocation that is also complete.

In this work, our objective is to obtain iterative auctions that allocate items to bidders according to an optimal solution of the above optimization problem. In the subsequent sections, we provide alternative (linear) optimization formulations for the efficient allocation problem, and discuss how we can use those to develop iterative auctions for graphical valuations.

4.2.2 Anonymous Item Pricing and Walrasian Equilibrium

In iterative auctions the auctioneer sets prices for the (subsets of) items she is offering, collects demand reports from bidders, and adjusts the prices. Before we discuss the details of the iterative auction design problem, a fundamental question to answer is the choice of the pricing rule for these auctions.

An important pricing rule that can be used for iterative auction design is the *anonymous pricing rule*, where the auctioneer offers the price p_i for each item $i \in \mathcal{N}$, which is the same for all bidders. Due to its simplicity, this pricing rule is commonly used in practice, and various theoretical works develop iterative auctions based on it (Ausubel, 2004, 2006; Gul and Stacchetti, 1999, 2000).

At given anonymous item prices, we say that a bundle S^* is *demanded* by bidder m , if maximum surplus is achieved for this bundle, i.e., $v^m(S^*) - \sum_{i \in S^*} p_i = \max_S v^m(S) - \sum_{i \in S} p_i$. We denote the set of bundles a bidder demands by D^m , i.e., $D^m = \arg \max_S v^m(S) - \sum_{i \in S} p_i$.

A natural termination condition for iterative auctions that rely on anonymous item pricing rule is “market clearance”. In particular, the auctioneer can terminate the auction if bidders demand disjoint sets of items, and every item is demanded by a bidder. Observe that in such an outcome, since demand sets are disjoint, bidders do not compete with the remaining bidders for the items that they demand. Hence, the auctioneer can terminate the auction by assigning each bidder the set of items she demands.

This outcome coincides with the well-known Walrasian equilibrium concept which we discuss next:

Definition 4.2.4 (Walrasian equilibrium). Let p_i denote the price of item i , and S^m denote the set of items assigned to player m . The tuple $(\{p_i\}_i, \{S^m\}_m)$ is a Walrasian equilibrium if:

- $p_1, \dots, p_N \geq 0$,
- S^1, \dots, S^M is a feasible allocation, i.e., $S^k \cap S^m = \emptyset$,
- $v^m(S^m) - \sum_{i \in S^m} p_i \geq v^m(S) - \sum_{i \in S} p_i$ for any $S \subset \mathcal{N}$,
- $p_i = 0$ if $i \notin \cup_m S^m$.

Observe that at a Walrasian equilibrium, S^m corresponds to a set of items bidder m demands. By Assumption 4.2.1, it follows that if a Walrasian equilibrium exists, then there exists one where $\{S^m\}$ is a complete feasible allocation. This suggests that in such an allocation, $\{S^m\}$ clears the market by assigning each bidder a set of items she demands. Additionally, it is known that if a Walrasian equilibrium exists the allocation $\{S^m\}$ is efficient (Gul and Stacchetti, 1999).

The above discussion and Definition 4.2.4 imply that iterative auctions that rely on anonymous item pricing can naturally be terminated when a Walrasian equilibrium is reached. Moreover, simple price update processes (such as tâtonnement) converge to a Walrasian equilibrium and the efficient allocation. This observation was used in the existing literature (Ausubel, 2006) for iterative auction design.

On the other hand, in an economy with indivisibilities a Walrasian equilibrium need not always exist. An important case, where it is known to exist, is the case of gross substitutes (Gul and Stacchetti, 1999). However, valuation functions satisfying the gross substitutes condition can only exhibit substitutability in valuations, and not complementarity.

A necessary and sufficient condition is present in the literature for testing the existence of Walrasian equilibrium. In particular, in Bikhchandani and Mamer (1997) authors establish that a Walrasian equilibrium exists if and only if the following linear program has optimal solutions that are integral:

$$\begin{aligned}
 \max \quad & \sum_m \sum_S x^m(S) v^m(S) \\
 \text{s.t.} \quad & \sum_S x^m(S) \leq 1 \quad \text{for all } m \\
 (LP1) \quad & \sum_m \sum_{S|i \in S} x^m(S) \leq 1 \quad \text{for all } i \\
 & x^m(S) \geq 0.
 \end{aligned} \tag{4.1}$$

At a feasible integral solution of this optimization problem, $x^m(S) = 1$ captures assignment of bundle S to bidder m . The first constraint suggests that each bidder receives at most one bundle. The second constraint, on the other hand, suggests that each item i can be

present in at most one bidder's bundle. The objective value is the welfare generated by the allocation suggested by $\{x^m(\cdot)\}_m$.

The corresponding dual program (D1) is stated below:

$$\begin{aligned}
 & \min \sum_i p_i + \sum_m \pi^m \\
 (D1) \quad & s.t. \pi^m \geq v^m(S) - \sum_{i \in S} p_i \quad \forall S, m \\
 & p_i, \pi^m \geq 0.
 \end{aligned}$$

Observe that in the dual LP, we have a variable p_i for each item i . This variable can be interpreted as the price the auctioneer offers for the relevant item. We also have a variable π^m for each bidder m . The constraints in the dual problem suggest that for any bundle S , π^m is an upper bound on $v^m(S) - \sum_{i \in S} p_i$, i.e., the surplus bidder m associates with acquiring bundle S at the given prices. Moreover, at optimality, it can be seen that π^m will be equal to the maximum surplus that can be associated with the bundle bidder m demands. This suggests that π^m can be interpreted as the surplus of bidder m . The results of Bikhchandani and Mamer (1997) imply that when a Walrasian equilibrium exists the $\{p_i\}$ variables at an optimal solution of D1, together with the allocation $\{S^m\}$ obtained at an optimal solution of LP1 (such that $x^m(S^m) = 1$ for all m) constitute a Walrasian equilibrium.

In Chapter 6, we develop iterative auctions by obtaining solutions of the above primal-dual LP pair (as well as a more general pair that we provide in Chapter 5) using iterative algorithms. These algorithms start with a dual feasible solution, and iteratively adjust this solution until optimal solutions to both LPs are found (hence a Walrasian equilibrium is identified). They can be interpreted as setting prices ($\{p_i\}$) for the items in the auction, and collecting bidders' demand at these prices, and adjusting the prices appropriately at each step of the auction. The corresponding auction formats rely on an *anonymous item pricing rule*, where the auctioneer offers a price (p_i) for each item i , and this price is the same for all bidders.

4.3 Walrasian Equilibrium and Tree Valuations

In the rest of this chapter, we assume that the value functions are not arbitrary graphical valuations, but they have additional structure. In particular, we assume that the underlying value graph is a tree graph, i.e., it contains no cycles. Additionally we assume that if two items are complements (substitutes) for a given player, they are complements (substitutes) for all players. In this setting, we establish that a Walrasian equilibrium exists. This result implies that LP1 of Section 4.2 can be used to identify the efficient allocation for sign-consistent tree valuations.

We start by formally stating the additional structure we impose on the valuation functions in the rest of the section:

Assumption 4.3.1 (Tree valuation). *Let $G = (\mathcal{N}, E)$ be the value graph associated with the graphical valuations of bidders. We assume that G is a tree graph, i.e., it contains no cycles.*

Assumption 4.3.2 (Sign Consistency). *The graphical valuations are sign-consistent. That is, for some $(i, j) \in E$ and $m \in \mathcal{M}$, if $w_{ij}^m > 0$, then $w_{ij}^k \geq 0$ for all $k \in \mathcal{M}$, and similarly if $w_{ij}^m < 0$, then $w_{ij}^k \leq 0$ for all $k \in \mathcal{M}$.*

The tree graph assumption, given in Assumption 4.3.1, imposes additional restrictions on the complementarity and substitutability structure. In particular, it disallows for complementarity and substitutability cycles. This structure may be reasonable in some auction settings. For instance, consider a spectrum auction, where one of the items correspond to a band in a central region, and the remaining correspond to bands in peripheral regions that are close to the central one. If the interference between the peripheral regions is negligible, and the interference between the central region and the remaining ones is relatively strong, then the corresponding value structure can naturally be represented as a tree graph.

The sign consistency requirement suggests that two items i and j are either substitutes or complements for all players. Note that this assumption still allows for presence of both complementary and substitutable items in the set of items, but it disallows having two items as substitutes for some players and complements for the remaining ones. More formally, the above assumption allows for positive and negative weights on different edges, but it requires that all players have same sign weight on a given edge. Observe that this assumption can also be satisfied by the spectrum auction example described above. For instance, if two bands are complementary for a player (due to reduced interference), we expect them to be complementary for all players.

We next show that under these assumptions a Walrasian equilibrium exists, or equivalently LP1 has optimal solutions that are integral and that correspond to the efficient allocation. In order to establish this result, in Section 4.3.1, we first consider a related LP formulation, which more clearly makes use of the structure of graphical valuations, and establish that this LP has integral optimal solutions for sign-consistent tree valuations. Then, in Section 4.3.2, we show that this result implies the existence of an integral solution of LP1 for the aforementioned class of valuation functions.

4.3.1 An Alternative LP Formulation

If the valuation functions of bidders were public knowledge, the efficient allocation could be found by solving an integer program. One possible integer programming formulation that can be used when the valuations of players are represented by graphical valuations is

provided below:

$$\begin{aligned}
\max \quad & \sum_m \sum_{i \in \mathcal{N}} x_i^m w_i^m + \sum_m \sum_{ij \in E} y_{ij}^m w_{ij}^m \\
\text{s.t.} \quad & \sum_m x_i^m \leq 1 && i \in \mathcal{N}, \\
(IP) \quad & y_{ij}^m \leq x_i^m, && ij \in E, \\
& y_{ij}^m \leq x_j^m, \\
& x_i^m + x_j^m - 1 \leq y_{ij}^m, && ij \in E, \\
& x_i^m \in \{0, 1\}, y_{ij}^m \in \{0, 1\}, && i \in \mathcal{N}, ij \in E
\end{aligned} \tag{4.2}$$

We use the shorthand notation (x, y) to denote a feasible solution of this problem given by $\{x_i^m\}_{m \in \mathcal{M}, i \in \mathcal{N}}, \{y_{ij}^m\}_{m \in \mathcal{M}, ij \in E}$.

In this formulation, $x_i^m = 1$ captures assigning item i to player m . We set $y_{ij}^m = 1$ if both item i and item j are assigned to player m , and the corresponding edge weight needs to be taken into account, when computing the total value of player m for her bundle. The first constraint guarantees that each item is assigned to at most one player. The second and third constraints guarantee that if $y_{ij}^m = 1$, then $x_i^m = x_j^m = 1$, i.e., if the weight of edge (i, j) is taken into account when computing the value of player m for her bundle, then it should be the case that this player receives both item i and item j . The fourth constraint states that if $x_i^m = x_j^m = 1$, then $y_{ij}^m = 1$, i.e., if player m receives both item i and item j then the value of edge (i, j) should be taken into account when computing the total value she has for her bundle. We next establish that the optimization formulation IP can be used to find the efficient allocation.

Lemma 4.3.1. (i) *Every feasible allocation $\{S^m\}_m$, corresponds to a unique feasible solution of IP, denoted by (x, y) , and vice versa. Moreover, (x, y) is such that $x_i^m = 1$ if $i \in S^m$, and $x_i^m = 0$ otherwise.*

(ii) *Additionally, the objective value of (IP) for the feasible solution (x, y) equals to $\sum_m v^m(S^m)$.*

(iii) *If (x, y) is an optimal solution of (IP), then the corresponding allocation $\{S^m\}_m$ is efficient.*

Proof. (i) Observe that given x (such that $x_i^m \in \{0, 1\}$ for all m, i), there is a unique y such that (x, y) is feasible in (IP). To see this focus on edge (i, j) and consider the following three cases: (a) $x_i^m = x_j^m = 0$, (b) $x_i^m = x_j^m = 1$, (c) $x_i^m = 0, x_j^m = 1$. In cases (a) and (c) the (second and third) constraints of (IP) indicate that $y_{ij}^m = 0$ in a feasible solution. Similarly, in case (b) the fourth constraint of (IP) indicates that $y_{ij}^m = 1$ in a feasible solution. Thus, in all cases, specifying x , uniquely identifies the corresponding y in a feasible solution.

Consider a feasible allocation $\{S^m\}_m$, and define x such that $x_i^m = 1$ if $i \in S^m$, and $x_i^m = 0$ otherwise. Clearly the above mapping between the set of feasible allocations $\{S^m\}_m$, and x (such that $x_i^m \in \{0, 1\}$ for all m, i) is a bijection. Since, specifying such an x uniquely identifies the corresponding y in the feasible solution, it follows that there is a bijection between the set of feasible allocations and the feasible solutions of (IP), and the claim follows.

- (ii) Using the definition of graphical valuations, the total value corresponding to a feasible allocation can be given as

$$\sum_m v^m(S^m) = \sum_m \left(\sum_{i \in S^m} w_i^m + \sum_{i,j \in S^m | (i,j) \in E} w_{ij}^m \right). \quad (4.3)$$

It follows from our construction in part (i) that $x_i^m = 1$ if and only if $i \in S^m$, in the solution (x, y) corresponding to $\{S^m\}$. This implies that $\sum_{i \in S^m} w_i^m = \sum_{i \in \mathcal{N}} x_i^m w_i^m$. Analogously, our construction suggests that $i, j \in S^m$ if and only if $x_i^m = x_j^m = 1$. Moreover, the constraints of (IP) imply that this is the case if and only if $y_{ij}^m = 1$ (assuming $(i, j) \in E$). Thus, $\sum_{i,j \in S^m | (i,j) \in E} w_{ij}^m = \sum_{(i,j) \in E} w_{ij}^m y_{ij}^m$. These observations together with (4.3) imply that the objective value of (IP) for the feasible solution (x, y) equals to $\sum_m v^m(S^m)$.

- (iii) Let (x, y) be an optimal solution of (IP), and $\{S^m\}$ be the corresponding feasible allocation. Consider any other feasible allocation $\{Z^m\}$, and the corresponding feasible solution of (IP). It follows from optimality of (x, y) that the solution corresponding to $\{Z^m\}$ leads to (weakly) lower objective value in (IP), when compared to (x, y) . On the other hand, by (ii) this implies that the corresponding feasible allocations are such that $\sum_m v^m(S^m) \geq \sum_m v^m(Z^m)$. Since this is true for any feasible allocation $\{Z^m\}$, it follows that $\{S^m\}$ is efficient. \square

We next focus on the LP relaxation of the problem (IP):

$$\begin{aligned} \max \quad & \sum_m \sum_{i \in \mathcal{N}} x_i^m w_i^m + \sum_m \sum_{ij \in E} y_{ij}^m w_{ij}^m \\ \text{s.t.} \quad & \sum_m x_i^m \leq 1 && i \in \mathcal{N}, \\ (LP2) \quad & y_{ij}^m \leq x_i^m, && ij \in E, \\ & y_{ij}^m \leq x_j^m, && ij \in E, \\ & x_i^m + x_j^m - 1 \leq y_{ij}^m, && ij \in E, \\ & 0 \leq x_i^m \leq 1, && i \in \mathcal{N}, \\ & 0 \leq y_{ij}^m \leq 1, && ij \in E. \end{aligned} \quad (4.4)$$

In general this LP relaxation is not exact, i.e., the relaxation may have non-integral solutions and lead to a higher objective value than the optimal objective value of (IP). We provide examples for this in Section 4.4. Interestingly, when Assumptions 4.3.1 and 4.3.2 hold, we can establish that this LP relaxation always has optimal solutions that are integral. This claim is formalized in the next theorem.

Theorem 4.3.1. *Assume that valuation functions of bidders are graphical and Assumptions 4.3.1 and 4.3.2 hold. Then, LP2 always has an optimal solution, which is also optimal for (IP).*

Proof. We prove the claim by first considering a relaxation of LP2, and showing that this problem has an optimal solution when Assumption 4.3.2 holds. Then, we show that the extreme points of the feasible region of this new problem are integral, and they are also feasible in LP2. The claim then follows from the linear structure of the problem.

Observe that by Assumption 4.3.2, it follows that for a given edge, either all players have positive weights, or all players have negative weights. Let $E^+ \subset E$ denote the subset of edges, for which player weights are nonnegative ($w_{ij}^m \geq 0$), and $E^- \subset E$ denote the subset, for which the weights are nonpositive ($w_{ij}^m \leq 0$).

We consider a new optimization problem, LP2b, that is obtained by relaxing some of the constraints in LP2:

$$\begin{aligned}
\max \quad & \sum_m \sum_{i \in \mathcal{N}} x_i^m w_i^m + \sum_m \sum_{ij \in E} y_{ij}^m w_{ij}^m \\
\text{s.t.} \quad & \sum_m x_i^m \leq 1 && i \in \mathcal{N}, \\
\text{(LP2b)} \quad & y_{ij}^m \leq x_i^m, && ij \in E^+, \\
& y_{ij}^m \leq x_j^m, && \\
& x_i^m + x_j^m - 1 \leq y_{ij}^m, && ij \in E^-, \\
& 0 \leq x_i^m, && i \in \mathcal{N}, \\
& 0 \leq y_{ij}^m && ij \in E^-.
\end{aligned} \tag{4.5}$$

Observe that LP2b is obtained from LP2 by relaxing (i) the constraint $y_{ij}^m \leq 1$ for all edges (i, j) , (ii) the constraints $x_i^m \leq 1$ for all nodes i , (iii) the upper bound constraints $y_{ij}^m \leq x_i^m, x_j^m$ for edges $(i, j) \in E^-$, and (iv) the lower bound constraints $x_i^m + x_j^m - 1 \leq y_{ij}^m$ for edges $(i, j) \in E^+$. It can be seen that the constraints (i) and (ii) can be omitted without changing the feasible region, as the constraints $x_i^m \geq 0$, $\sum_m x_i^m \leq 1$ and $y_{ij}^m \leq x_i^m$ imply these constraints. Omitting the constraints in (iii) and (iv), on the other hand, changes the feasible region. In particular, it can be seen that in LP2b, the feasible region is unbounded, whereas the feasible region of LP2 is bounded. For instance, it can be seen that in LP2b we can have $y_{ij}^m \gg 1$ for $(i, j) \in E^-$ at a feasible solution.

We next establish that LP2b always has an optimal solution (see Section 4.7 for a proof).

Lemma 4.3.2. *Let Assumption 4.3.2 hold. Then LP2b has an optimal solution.*

If LP2b has an optimal solution, it should be at one of the extreme points of its feasible region. We next establish that the extreme points of this problem are all integral. The following lemma, proof of which can be found in Section 4.7, allows us to characterize the extreme points of the feasible region of LP2b:

Lemma 4.3.3. *Let Assumptions 4.3.1 and 4.3.2 hold. Then, all extreme points of LP2b are such that $x_i^m \in \{0, 1\}$ for all m and i .*

It also follows from this lemma that at an extreme point of LP2b, $y_{ij}^m = \min\{x_i^m, x_j^m\} \in \{0, 1\}$ for $ij \in E^+$. Assume that this is not the case, then no constraint that involves y_{ij}^m is binding and we obtain other feasible solutions by just considering $y_{ij}^m + \epsilon$, $y_{ij}^m - \epsilon$ instead of y_{ij}^m , and keeping the remaining elements of the feasible solution (x, y) intact. Thus, we obtain a contradiction and at an extreme point $y_{ij}^m = \min\{x_i^m, x_j^m\} \in \{0, 1\}$ for $ij \in E^+$. Similarly, it follows that $y_{ij}^m = \max\{x_i^m + x_j^m - 1, 0\} \in \{0, 1\}$ for $ij \in E^-$ at an extreme point of LP2b.

We next show that any extreme point of LP2b is feasible in LP2 when Assumptions 4.3.1 and 4.3.2 hold. To establish this, note that it is sufficient to check that the extreme points satisfy the relaxed constraints (i)-(iv) of LP2 (stated after (4.5)), as the remaining constraints are already satisfied by feasibility in LP2b. As explained above, when the assumptions hold, the extreme points are such that (a) $x_i^m \in \{0, 1\}$, (b) $y_{ij}^m = \max\{x_i^m + x_j^m - 1, 0\} \in \{0, 1\}$ for $ij \in E^-$, and (c) $y_{ij}^m = \min\{x_i^m, x_j^m\} \in \{0, 1\}$ for $ij \in E^+$, is feasible for LP2. Any such point immediately satisfies $y_{ij}^m \leq 1$ for all edges (i, j) , and $x_i^m \leq 1$ for all nodes i (i.e., constraints (i) and (ii)). The upper bound constraints $y_{ij}^m \leq x_i^m, x_j^m$ for edges $(i, j) \in E^-$ (i.e., constraint (iii)), hold since $y_{ij}^m = 1$ only when $x_i^m = x_j^m = 1$ (recall that for $(i, j) \in E^-$, we have $y_{ij}^m = \max\{x_i^m + x_j^m - 1, 0\}$). The lower bound constraints $x_i^m + x_j^m - 1 \leq y_{ij}^m$ for edges $(i, j) \in E^+$ (i.e., constraint (iv)), hold since $y_{ij}^m = 0$ only when $x_i^m = 0$ or $x_j^m = 0$ (recall that for $(i, j) \in E^+$, we have $y_{ij}^m = \min\{x_i^m, x_j^m\}$).

Summarizing, using Lemma 4.3.3 we conclude that when Assumptions 4.3.1 and 4.3.2 hold, any extreme point of LP2b is a feasible integral solution of LP2. In addition, when Assumption 4.3.2 holds, an optimal solution to LP2b exists and is one of these extreme points (Lemma 4.3.2). Hence, we conclude that when both assumptions hold the optimal solution of LP2b is a feasible integer solution of LP2. Since LP2b is obtained by relaxing some constraints of LP2, this solution is also optimal in LP2. Hence, we conclude that LP2 has an optimal integer solution under Assumptions 4.3.1 and 4.3.2. \square

This result suggests that when Assumptions 4.3.1 and 4.3.2 hold, the efficient allocation can be found by solving the linear optimization problem LP2. In Section 4.2.2, we establish that this result also implies that LP1 has integral optimal solutions under these assumptions.

4.3.2 Existence of a Walrasian Equilibrium

We next establish a relation between LP1 and LP2, and subsequently use this to conclude that a Walrasian equilibrium exists, when Assumptions 4.3.1 and 4.3.2 hold.

Theorem 4.3.2. *If LP2 has an optimal solution that is integral, then so does LP1.*

The proof of this theorem can be found in Section 4.7. The main idea behind the proof is to establish that for any feasible solution of LP1, it is possible to construct a feasible solution of LP2 with the same objective value, and conversely for any feasible integral solution of LP2, it is possible to construct a feasible integral solution of LP1 again with the same objective value (see Figure 4-3). These two facts immediately imply that when LP2 has an optimal solution that is integral, this solution leads to a (weakly) larger objective value than all feasible solutions of LP1. Moreover, there exists a feasible integral solution of LP1 with the same objective value. Thus, this solution is an optimal solution of LP1.




Figure 4-3: A feasible solution $\{x_i^m, y_{ij}^m\}$ to LP2 can be constructed from a feasible solution $\{x^m(S)\}$ of LP1 (by setting $x_i^m = \sum_{S|i \in S} x^m(S)$, $y_{ij}^m = \sum_{S|ij \in S} x^m(S)$). These solutions have the same objective values in the corresponding optimization problems. Additionally, the feasible integer points of LP2 correspond to the feasible integer points of LP1. Thus, if LP2 admits an optimal solution that is integral, then so does LP1.

An immediate consequence of this result is that a Walrasian equilibrium exists for tree valuations that satisfy the sign consistency condition.

Corollary 4.3.1. *Let Assumptions 4.3.1 and 4.3.2 hold. Then, LP1 has an integral optimal solution and a Walrasian equilibrium exists.*

Proof. It was established in Theorem 4.3.1 that under these assumptions LP2 has an integral optimal solution. Theorem 4.3.2, implies that in this case LP1 also has an integral optimal solution. However, as stated above this is a necessary and sufficient condition for existence of a Walrasian equilibrium. \square

Since graphical valuations can capture both value complementarity and substitutability, this result implies that a Walrasian equilibrium exists in settings where both are present, provided that the value structure satisfies Assumption 4.3.1 and Assumption 4.3.2.³ As we demonstrate in Chapter 6, this result can be used to obtain iterative auction formats that rely on anonymous item pricing, terminate at a Walrasian equilibrium, and guarantee efficiency for sign-consistent tree valuations.

4.4 Relaxing the Assumptions

In the previous section, under Assumptions 4.3.1 and 4.3.2, we provided LP formulations that can be used to find the efficient outcome, and established the existence of Walrasian equilibrium. In this section, we show that if we relax these assumptions, then a Walrasian equilibrium need not exist, and the LP formulations LP1 and LP2 need not have optimal solutions that are integral.

4.4.1 Sign Consistency

We first focus on the sign consistency condition of Assumption 4.3.2, and show that if players have edge weights with different signs, the LP formulations do not lead to integral solutions.

Consider a setting with two bidders m, k and two items i, j . Assume that the valuations of bidders are represented with graphical valuations (see Figure 4-4), such that

- $w_i^m = w_j^m = w_{ij}^m = 5$.
- $w_i^k = w_j^k = 10, w_{ij}^k = -10$.

Observe that the graphical valuation in this example satisfies Assumption 4.3.1 but not Assumption 4.3.2.




Figure 4-4: The weights for player m is given above the nodes/edges, and those for player k are given below them.

Note that the optimal integral solutions of LP1 result in a total welfare of 15 (this can be obtained either by assigning both items to player m , i.e., $x^m(\{i, j\}) = x^k(\emptyset) = 1$, and

³Another class of valuations for which a Walrasian equilibrium exists, and valuations can exhibit both complementarity and substitutability, is the class of gross substitutes and complements (Sun and Yang, 2006). The class of valuations that satisfy Assumptions 4.3.1 and 4.3.2, on the other hand, is not contained in this class, as explained in Section 4.6. Thus, our result here establishes existence of a Walrasian equilibrium for a distinct and important class of valuations.

$x^l(S) = 0$, for remaining S and $l \in \{m, k\}$; or assigning one item to m and the other one to k , i.e., $x^m(\{i\}) = x^k(\{j\}) = 1$ and $x^l(S) = 0$, for remaining S and $l \in \{m, k\}$. On the other hand, consider the following solution of LP1: $x^m(\{i, j\}) = x^m(\emptyset) = 1/2$, $x^k(\{i\}) = x^k(\{j\}) = 1/2$, and $x^m(\{i\}) = x^m(\{j\}) = x^k(\{i, j\}) = x^k(\emptyset) = 0$. Feasibility of this solution can be immediately checked. The objective value associated with this solution equals to

$$\begin{aligned} x^m(\{i, j\})v^m(\{i, j\}) + x^m(\emptyset)v^m(\emptyset) + x^k(\{i\})v^k(\{i\}) + x^k(\{j\})v^k(\{j\}) \\ = \frac{1}{2}(15 + 0 + 10 + 10) = 17.5. \end{aligned}$$

Note that the objective value associated with this solution is larger than the objective of the optimal integral solution of LP1. Thus, LP1 does not have an optimal solution that is integral. This implies that a Walrasian equilibrium does not exist for the given valuations. Additionally, by Theorem 4.3.2, we conclude that for this example LP2 cannot have an optimal solution that is integral.

Observe that the fact that the LP formulations do not have optimal integral solutions imply that the underlying polytopes are not integral, and they have extreme points with nonintegral coordinates. Remarkably, despite presence of such extreme points, Theorem 4.3.1 and Corollary 4.3.1 establish that LP2 and LP1 have optimal solutions that are integral when Assumptions 4.3.1 and 4.3.2 hold (observe that Assumption 4.3.2 has no impact on the underlying feasible region, hence extreme points with nonintegral coordinates exist even when this assumption holds).

4.4.2 Value Graphs with Cycles

Next, we establish that if the underlying graph has a k -cycle (i.e., a cycle that involves k nodes, where $k > 2$), then the LP formulations can have optimal solutions that are not integral.

Assume that there are N nodes (items), and the nodes $1, \dots, k$ of the underlying graph are in the form of a cycle.

We first consider the case where k is odd. Let there be $M = k$ players. Further assume that all players have weights equal to zero for all nodes. We associate nonzero edge weights only with the edges contained in the cycle, i.e., edges (i, j) such that $1 \leq i, j \leq k$, $j = i + 1$ in mod k . In particular,

- For players $m \in \{1, \dots, k - 1\}$, we assume that $w_{ij}^m = \epsilon > 0$ if $i = m$, $j = m + 1$ and $w_{ij}^m = 0$ otherwise.
- For player $m = k$, we assume that $w_{ij}^m = \epsilon > 0$ if $i = m$, $j = 1$ and $w_{ij}^m = 0$ otherwise.

Observe that the value of the optimal integral solution of LP1 is upper bounded by $\epsilon \lfloor \frac{k}{2} \rfloor$. This is because if the weight of edge $(i, i + 1)$ contributes to the objective by ϵ , it

should be that player i (who is the only player with positive value for this edge) receives the items at its end points. However, this implies that the weights of the neighboring edges $(i-1, i)$ and $(i, i+1)$ do not contribute to the objective (since player i has value equal to 0 for these edges, but she acquires nodes i and $i+1$, no other player can benefit from the edge value of these edges). Thus, at most half of the edges can contribute to the welfare, and the resulting welfare is bounded by $\epsilon \lfloor \frac{k}{2} \rfloor$ (where floor operation is present, since at an integral solution the total welfare will always be an integer multiple of ϵ).

On the other hand, consider the following solution of LP1:

- For players $m \in \{1, \dots, k-1\}$ we assume that $x^m(\{m, m+1\}) = x^m(\emptyset) = 1/2$ and $x^m(S) = 0$ for the remaining bundles S .
- For player $m = k$, we assume that $x^m(\{m, 1\}) = x^m(\emptyset) = 1/2$ and $x^m(S) = 0$ for the remaining bundles S .

The feasibility of this solution in LP1 can immediately be checked. Observe that the objective value of LP1 associated with this solution is

$$\sum_m \sum_S x^m(S) v^m(S) = \frac{1}{2} \sum_m \epsilon = \epsilon k/2.$$

Thus, when k is odd, we observe that the optimal value of LP1 is larger than the value of the optimal integral solution. Thus, LP1 does not have an optimal solution that is integral. This implies that (i) LP2 does not have an optimal solution that is integral (from Theorem 4.3.2), (ii) a Walrasian equilibrium does not exist.

If k is even, we slightly modify the above construction. In this case, we assume that there are $m = k-1$ players. As before, players have zero weights for all nodes, and the edge weights are such that

- For players $m \in \{1, \dots, k-2\}$, we assume that $w_{ij}^m = \epsilon > 0$ if $i = m, j = m+1$ or $i = m+1, j = m+2$ and $w_{ij}^m = 0$ otherwise.
- For player $m = k-1$, we assume that $w_{ij}^m = \epsilon > 0$ if $i = k, j = 1$ and $w_{ij}^m = 0$ otherwise.

It can be shown that the value of the objective at an integral solution of LP1 is upper bounded by $\epsilon(k-2)$. In particular, there are two cases to consider: (i) player $k-1$ receives a bundle containing edge $(k, 1)$, (ii) player $k-1$ does not receive a bundle containing this edge. In case (i), it can be seen that the contribution of edges $(1, 2)$ and $(k-1, k)$ to the objective function is equal to zero (since at least one end point of these edges is assigned to player k who has no value for the corresponding edges). Consequently, the maximum welfare is bounded by $\epsilon(k-2)$. In case (ii), it can be seen that the total welfare is immediately bounded by $\epsilon(k-1)$, since none of the other players other than $k-1$ have a positive value

for edge $(k, 1)$. On the other hand, for total welfare to be equal to $\epsilon(k - 1)$, we need all remaining edges to have a contribution of ϵ to the objective. This requires assigning all edges to a single bidder (otherwise, weights associated with some edges do not contribute to the objective, since the end points of some edges are not assigned to a single bidder). However, no single bidder has a strictly positive weight for more than 2 edges. Thus, it follows that in this case the welfare is strictly less than $\epsilon(k - 1)$. Since for integral solutions, welfare is a multiple of ϵ , it follows that in case (ii) the total welfare is bounded by $\epsilon(k - 2)$.

We consider the following feasible solution of LP1:

- For players $m \in \{1, \dots, k - 2\}$ we assume that $x^m(\{m, m + 1, m + 2\}) = x^m(\emptyset) = 1/2$ and $x^m(S) = 0$ for the remaining bundles S .
- For player $m = k - 1$, we assume that $x^m(\{k, 1\}) = x^m(\emptyset) = 1/2$ and $x^m(S) = 0$ for the remaining bundles S .

It can be checked that this solution is feasible in LP1. Moreover, the corresponding objective value is given by

$$\sum_m \sum_S x^m(S) v^m(S) = \frac{1}{2} \sum_{m=1}^{k-2} 2\epsilon + \frac{1}{2}(\epsilon) = \epsilon \left(k - 2 + \frac{1}{2} \right).$$

As before, the optimal solution of LP1 leads to a larger objective value than the optimal integral solution. Hence, LP1 does not have an optimal solution that is integral. Thus, a Walrasian equilibrium does not exist, and LP2 cannot have an optimal solution that is integral.

For $k = 3$ and $k = 4$, the above constructions are illustrated in Figures 4-5 and 4-6. To simplify the figures, the nodes $1, 2, \dots$ in our construction are relabeled as A, B, \dots in the figures.

We conclude that both the tree graph assumption, and the sign consistency condition are critical for the existence of a Walrasian equilibrium, and the existence of optimal solutions to LP formulations LP1 and LP2 that are integral. Hence, when these assumptions are relaxed, it may not be possible to implement the efficient outcome using iterative auction formats that rely on anonymous item pricing rule, and terminate when a market clearance condition holds. In the next chapter, we establish that in those cases efficient iterative auctions can still be developed by considering more general pricing rules.

4.5 Summary

In this chapter, we focused on a special class of graphical valuations, where the underlying value graph is a tree, and edge weights satisfy a sign consistency condition. We established that under these assumptions, a Walrasian equilibrium always exists. Additionally, the existence of a Walrasian equilibrium immediately suggests a linear programming formulation




Figure 4-5: Value graphs for k -cycles (left $k = 3$, right $k = 4$). We assume there are three players, and node weights are equal to zero for all players. The labels associated with the edges designate the edge weights for players (scaled by $1/\epsilon$). In particular, if the edge label is the triple w_1, w_2, w_3 , this implies that each player i has weight $w_i\epsilon$ associated with edge i . For the figure on the left, the optimal integral solution of LP1 results in an objective value of ϵ , whereas for the figure on the right, this value is 2ϵ .




Figure 4-6: For the valuations given in Figure 4-5, the optimal solution of LP1 is not integral. It can be seen that the fractional assignments given above lead to larger objective values than the integral solutions provided in Figure 4-5. The solution given on the left suggests setting $x^1(\{AB\}) = x^2(\{BC\}) = x^3(\{CA\}) = 1/2$ and $x^1(\emptyset) = x^2(\emptyset) = x^3(\emptyset) = 1/2$. The solution on the right suggests, setting $x^1(\{ABC\}) = x^2(\{BCD\}) = x^3(\{DA\}) = 1/2$ and $x^1(\emptyset) = x^2(\emptyset) = x^3(\emptyset) = 1/2$. The objective value corresponding to the solution on the left is $\frac{3}{2}\epsilon$, the optimal value corresponding to the solution on the right is $\frac{5}{2}\epsilon$. Observe that in both cases the obtained solutions result in strictly larger objective value than the corresponding optimal integer solution.

that can be used to identify the efficient allocation. We also demonstrated that if the tree, or the sign consistency assumption is relaxed, then the LP formulations that we provide in this section may not allow for finding the efficient outcome, and a Walrasian equilibrium need not exist. This suggests that in order to identify the efficient outcome for more general graphical valuations, a different and more complex LP formulation may be necessary. We provide such LP formulations in Chapter 5.

4.6 Appendix: Gross Substitutes and Complements Condition and Graphical Valuations

In this section, we explain how graphical valuations and tree valuations are different from other well-studied classes of valuation functions in the literature. In particular, we focus on the classes of gross substitutes and complements, gross substitutes, sub/superadditive, sub/supermodular valuation functions, and compare those with graphical valuations.

It was established in Section 4.3 that when the underlying value graph has a tree structure, and the valuations satisfy sign consistency, a Walrasian equilibrium exists. This result allows us to identify a class of valuation functions which exhibit both value complementarity and substitutability, and for which a Walrasian equilibrium exists. Gross substitutes and complements (Sun and Yang, 2006, 2009), defined below, is another class of valuation functions that satisfies a similar property.

Definition 4.6.1 (Gross Substitutes and Complements (GSC)). *Assume that the set of items is partitioned into two sets S_1, S_2 such that $S_1 \cap S_2 = \emptyset$, $S_1 \cup S_2 = \mathcal{N}$. Consider the valuation function $v : \mathcal{N} \rightarrow \mathbb{R}$. Denote by $e(k)$ the k th unit vector, and $D(p)$ the demand function associated with price vector $p \in \mathbb{R}^N$, i.e., $D(p) \triangleq \arg \max_{S \subset \mathcal{N}} v(S) - \sum_{i \in S} p_i$.*

We say that v has the gross substitutes and complements property if for $j \in \{1, 2\}$, any price vector $p \in \mathbb{R}^N$, $k \in S_j$, $\delta \geq 0$, and $D_1 \in D(p)$, there exists $D_2 \in D(p + \delta e(k))$ such that (a) $[D_1 \cap S_j] - \{k\} \subset D_2$ and (b) $D_1^c \cap S_j^c \subset D_2^c$.

Intuitively, this definition suggests that the items in sets S_1 and S_2 are substitutes among themselves (in the sense that if the price of a demanded item in one of these sets increases, the demand for the other demanded items in the same set does not decrease, $[D_1 \cap S_j] - \{k\} \subset D_2$). Additionally items are complements across S_1 and S_2 (in the sense that if the price of a demanded item in set S_1 increases, then fewer items are demanded in set S_2 , $D_1^c \cap S_j^c \subset D_2^c$).

We next illustrate that tree (and hence graphical) valuations need not satisfy the GSC property, by considering the tree valuation provided in Figure 4-7. Assume that this valuation satisfies the GSC property. There are three different ways of choosing sets S_1 and S_2 (due to symmetry all other cases follow from the analysis here): (i) $S_1 = \{A, B, C\}$, $S_2 = \emptyset$, (ii) $S_1 = \{A\}$, $S_2 = \{B, C\}$, (iii) $S_1 = \{A, C\}$, $S_2 = \{B\}$.




Figure 4-7: A tree valuation that violates the GSC property.

We will show that the GSC property fails in all of these cases, and hence the value function given in Figure 4-7 does not exhibit the GSC property for any choice of $\{S_j\}$. Assume that

- Initially, the prices are $p_1(A) = 0.1$, $p_1(B) = 0.5$, $p_1(C) = 0.1$, and the corresponding demand is $D(p_1) = \{A, C\}$.
- Then, the price of the first item is increased, and the new prices are $p_2(A) = 1$, $p_2(B) = 0.5$, $p_2(C) = 0.1$. It follows that the demand is $D(p_2) = \{B\}$.

This implies that the GSC property fails whenever A and C belong to the same S_j (note that by choosing $D_1 = \{A, C\}$, $D_2 = \{B\}$, $k = A$, the condition $[D_1 \cap S_j] - \{k\} \subset D_2$ fails). Thus, to check the GSC property it is sufficient to focus on case (ii). On the other hand, if $S_1 = \{A\}$, $S_2 = \{B, C\}$ then the condition $D_1^c \cap S_j^c \subset D_2^c$ fails (this can be seen by choosing $j = 1$, $D_1 = \{A, C\}$, $D_2 = \{B\}$, $k = A$). This implies that the GSC property fails in case (ii) as well.

Hence, we conclude that for any choice of the $\{S_j\}$ sets, the GSC property fails for the value function in Figure 4-7. Thus, it follows that tree valuations are not contained in the class of GSC valuations. GSC generalizes the well-known gross substitutes class (Gul and Stacchetti, 1999), where Definition 4.6.1 holds with $S_2 = \emptyset$. Thus, our results also imply that tree valuations do not satisfy the gross substitutes property.

We emphasize that this conclusion still holds, if edge weights are not restricted to be negative as in the example, and allowed to be positive or negative. For instance, consider the tree valuation provided in Figure 4-8. Observe that this graph is obtained from Figure 4-7 after relabeling the nodes, and setting a positive edge weight to edge (A, C) . Here, the nodes are relabeled in order to have the same demand sets as in the previous example (Figure 4-7).




Figure 4-8: A tree valuation that violates the GSC property, and has edge weights with mixed signs.

Similar to the previous example assume that

- Initially, the prices are $p_1(A) = 0.1$, $p_1(C) = 1.5$, $p_1(B) = 0.1$, and the corresponding demand is $D(p_1) = \{A, C\}$.
- Then, the price of the first item is increased, and the new prices are $p_2(A) = 4$, $p_2(C) = 1.5$, $p_2(B) = 0.1$. It follows that the demand is $D(p_2) = \{B\}$.

Since demand sets are identical to those in the previous example (and only the price of the same item is increased), using similar arguments as before it follows that GSC condition does not hold for the example in Figure 4-8.⁴

We next investigate the additional structural assumptions under which graphical valuations exhibit the GSC property. Assume that the underlying value graph consists of connected components of size at most two. Note that in this case, valuations are additive over different connected components, and hence in order to test the GSC condition it suffices to restrict attention to subsets of the demand set that are contained in a given connected component of the graph.

Consider a pair of nodes (i, j) connected with an edge. Assume that $w_{ij}^m \leq 0$ and at a given price vector p item j belongs to a demand set D_1 . We claim that if the price of item i increases, then the demand for item j cannot decrease. The claim is immediate if $i \notin D_1$, i.e., i is not demanded at the original prices. Assume that $i \in D_1$. Observe that this implies that $w_j^m + w_{ij}^m - p_j \geq 0$, since otherwise bidder m can improve her payoff by not receiving item j at the price vector p , and hence $j \notin D_1$. On the other hand, since $w_{ij}^m \leq 0$, it follows that $w_j^m - p_j \geq w_j^m + w_{ij}^m - p_j \geq 0$. Thus, at the updated prices bidder m still maximizes her surplus by either receiving item j together with i or in isolation. Hence, item j belongs to a demand set after the price update, and condition (a) of Definition 4.6.1 holds by assigning items (i, j) to the same set S_1 or S_2 .

Conversely, assume that $w_{ij}^m \geq 0$ and at price vector p , item j does not belong to a demand set D_1 . We claim that if the price of item i increases, then the demand for item j cannot increase. As before, the claim is immediate if $i \notin D_1$. If $i \in D_1$, and $j \notin D_1$, then it should be the case that $w_j^m + w_{ij}^m - p_j \leq 0$. Moreover, since $w_{ij}^m \geq 0$, this implies that $w_j^m - p_j \leq w_j^m + w_{ij}^m - p_j \leq 0$. Note that after the price update this inequality continues to hold. Thus, it should be the case that there is a demand set to which item j does not belong after the price update. Hence, condition (b) of Definition 4.6.1 holds, by assigning items (i, j) to different sets S_1 and S_2 .

These observations imply that if the underlying graph consists of components of size at most two, the GSC condition holds, by assigning items that are connected with a positive weight to different sets S_1 and S_2 (see Definition 4.6.1), and items that are connected with

⁴Note that in this example, unlike the one in Figure 4-7, the case $S_1 = \{A, B\}$, $S_2 = \{C\}$ needs to be handled separately, as the edge weights are no longer symmetric. However, it immediately follows that in this case the GSC condition cannot hold. To see this consider increasing the price of item C (significantly), as opposed to A , in the example. After the price update the demand set becomes $\{A, B\}$, violating the GSC condition associated with this choice of the sets S_1, S_2 .

a negative edge to the same set.⁵ Since our examples indicate that even for trees with three nodes the GSC property can fail, this result suggests that only a restrictive subclass of tree valuations (that consist of connected components of size at most two) satisfies the GSC condition.

We close this section by discussing the relation of graphical valuations to subadditive/superadditive and submodular/supermodular valuations. A valuation function is subadditive if for any sets $A, B \subset \mathcal{N}$, it satisfies $v(A \cup B) \leq v(A) + v(B)$, and superadditive, if for disjoint A, B , it satisfies $v(A \cup B) \geq v(A) + v(B)$. Similarly a valuation function is submodular if for any sets A, B it satisfies $v(A \cup B) + v(A \cap B) \leq v(A) + v(B)$, and supermodular if $v(A \cup B) + v(A \cap B) \geq v(A) + v(B)$. These inequalities imply that for nonnegative valuation functions, submodularity implies subadditivity.

It can be easily checked that if all edge weights are positive (negative), graphical valuations are superadditive and supermodular (subadditive and submodular). On the other hand, if there is an edge with negative (positive) weight, the supermodularity/superadditivity (submodularity/subadditivity) condition cannot hold (consider A, B as singletons corresponding to the end points of this edge). Since the weights of different edges in our model can be positive or negative, it follows that even for the case of trees, graphical valuations are not contained in these classes. Hence, we conclude that the results presented in this chapter do not immediately follow from the known results for the aforementioned classes of valuation functions.

4.7 Appendix: Proofs and Additional Results

Proof of Lemma 4.3.2: Note that due to the linear structure of the problem (and nonemptiness of the feasible region), it is sufficient to prove that the objective value of this problem is bounded, to establish that the problem has an optimal solution.

Consider any feasible solution (x, y) of LP2b. Observe that $\sum_m x_i^m \leq 1$ and $x_i^m \geq 0$ implies that $x_i^m \leq 1$. Since $w_{ij}^m \leq 0$ for $(i, j) \in E^-$, $w_{ij}^m \geq 0$ for $(i, j) \in E^+$; and $y_{ij}^m \geq 0$ for $(i, j) \in E^-$, $y_{ij}^m \leq x_i^m, x_j^m \leq 1$ for $(i, j) \in E^+$ it follows that the objective value for (x, y) satisfies:

$$\begin{aligned} \sum_m \sum_{i \in N} x_i^m w_i^m + \sum_m \sum_{ij \in E} y_{ij}^m w_{ij}^m &= \sum_m \sum_{i \in N} x_i^m w_i^m + \sum_m \sum_{ij \in E^+} y_{ij}^m w_{ij}^m + \sum_m \sum_{ij \in E^-} y_{ij}^m w_{ij}^m \\ &\leq \sum_m \sum_{i \in N} w_i^m + \sum_m \sum_{ij \in E^+} w_{ij}^m. \end{aligned} \tag{4.6}$$

Thus, the objective value is bounded for any feasible solution of LP2b, and it follows that this problem has an optimal solution. \square

⁵A similar conclusion holds for gross substitute valuations: a graphical valuation satisfies the gross substitutes condition, if the graph consists of connected components of size at most two, and the edge weights are negative.

Proof of Lemma 4.3.3: We make use the following claim, to establish the lemma.

Claim 1. *Assume that there are at least two bidders, and Assumptions 4.3.1 and 4.3.2 hold. Let (x, y) be a feasible solution of (LP2b) such that for some node i , and player m , $x_i^m \in (0, 1)$; and $\sum_m x_j^m = 1$ for all j .*

For any player $k \neq m$ such that $x_i^k > 0$, and sufficiently small $\epsilon > 0$, there exist two feasible solutions (\hat{x}, \hat{y}) and (\bar{x}, \bar{y}) to (LP2b) such that:

- $\hat{x}_i^m = x_i^m + \epsilon$, $\hat{x}_i^k = x_i^k - \epsilon$
- $\bar{x}_i^m = x_i^m - \epsilon$, $\bar{x}_i^k = x_i^k + \epsilon$
- $\hat{x}_i^l = \bar{x}_i^l = x_i^l$ for $l \neq m, k$
- $(x, y) = \frac{1}{2}(\bar{x}, \bar{y}) + \frac{1}{2}(\hat{x}, \hat{y})$.

We will prove this claim, but first we show how this claim implies the lemma.

Observe that this claim immediately implies that the feasible solutions of LP2b for which (i) $x_i^{m^*} \in (0, 1)$ for some node i , and player m^* , (ii) $\sum_{m \in \mathcal{M}} x_j^m = 1$ for all j , cannot be an extreme point. This is because, at any such solution there are at least two bidders (for the second condition to hold at node i), and by the claim this solution can be expressed as a convex combination of two other feasible solutions (\bar{x}, \bar{y}) and (\hat{x}, \hat{y}) .

Consider an instance of LP2b with a set of agents \mathcal{M} , and a feasible solution of this problem for which (i) $x_i^{m^*} \in (0, 1)$ for some node i , and player m^* , (ii) and $\sum_{m \in \mathcal{M}} x_j^m < 1$ for some j . We next construct a relevant problem instance, and applying the claim to this problem instance, establish that the original solution to LP2b, denoted by (x, y) , cannot be an extreme point.

For each node j for which $\sum_{m \in \mathcal{M}} x_j^m < 1$, we define a fictitious agent s_j , and consider LP2b with the addition of these new agents. We define a new solution for this problem (z, ζ) such that $z_j^m = x_j^m$, $\zeta_{jk}^m = y_{jk}^m$ for all nodes j, k and agents $m \in \mathcal{M}$ present in the original problem. For each fictitious agent s_j we set $x_j^{s_j} = 1 - \sum_{m \in \mathcal{M}} x_j^m$, and $x_{i'}^{s_j} = 0$ for nodes $i' \neq j$, $\zeta_{i'j'}^{s_j} = 0$ for all edges $(i', j') \in E$. Observe that this solution is feasible in the new instance of LP2b with fictitious agents. This is because, the new solution is identical to the original solution for all agents but the fictitious ones. The variables corresponding to fictitious agents, on the other hand, are set such that sum of z_i^m variables over all players m ($m \in \mathcal{M}$, or m is a fictitious agent) is equal to 1 at all nodes. Since every fictitious agent m has nonzero z_i^m exactly at one node i (by construction), it follows that setting ζ equal to zero for all edges of fictitious agents does not violate any constraints associated with edges (i.e., constraints 2–4 in LP2b). These facts imply feasibility of (z, ζ) in the new formulation with fictitious agents.

Since $x_i^{m^*} \in (0, 1)$ in the original formulation, it follows that $z_i^{m^*} \in (0, 1)$ in our new formulation. Note that Claim 1 implies that there are feasible solutions $(\bar{z}, \bar{\zeta})$ and $(\hat{z}, \hat{\zeta})$ to LP2b with fictitious agents, such that $(z, \zeta) = \frac{1}{2}(\bar{z}, \bar{\zeta}) + \frac{1}{2}(\hat{z}, \hat{\zeta})$, such that $\bar{z}_i^{m^*} = z_i^{m^*} - \epsilon =$

$x_i^{m^*} - \epsilon$ and $\hat{z}_i^{m^*} = z_i^{m^*} + \epsilon = x_i^{m^*} + \epsilon$. Let (\bar{x}, \bar{y}) and (\hat{x}, \hat{y}) , respectively denote the restriction of $(\bar{z}, \bar{\zeta})$ and $(\hat{z}, \hat{\zeta})$ to agents other than the fictitious agents (i.e., \bar{x}, \hat{x} , correspond to the component of \bar{z}, \hat{z} associated with bidders $m \in \mathcal{M}$, and similarly for \bar{y}, \hat{y}).

Observe that these solutions are feasible in LP2b with set of agents \mathcal{M} . To show this, we only need to establish feasibility of constraints $\sum_{m \in \mathcal{M}} \bar{x}_j^m \leq 1$ and $\sum_{m \in \mathcal{M}} \hat{x}_j^m \leq 1$ for every node j , since feasibility of the remaining constraints follow from feasibility of $(\bar{z}, \bar{\zeta})$ and $(\hat{z}, \hat{\zeta})$ in the formulation with fictitious agents. On the other hand, the feasibility of $\sum_{m \in \mathcal{M}} \bar{x}_j^m \leq 1$ when the set of agents is \mathcal{M} follows because feasibility of $(\bar{z}, \bar{\zeta})$ in the formulation with fictitious agents implies that

$$\sum_{m \in \mathcal{M}} \bar{z}_j^m + \bar{z}_j^{s_j} = \sum_{m \in \mathcal{M}} \bar{x}_j^m + z_j^{s_j} = 1,$$

and hence $\sum_{m \in \mathcal{M}} \bar{x}_j^m = 1 - z_j^{s_j} \leq 1$. Similarly, feasibility of $\sum_{m \in \mathcal{M}} \hat{x}_j^m \leq 1$ follows from the feasibility of $(\hat{z}, \hat{\zeta})$ in the formulation with fictitious agents.

Thus, we conclude that (\bar{x}, \bar{y}) and (\hat{x}, \hat{y}) are feasible solutions of LP2b when the set of agents is equal to \mathcal{M} . Since, $(z, \zeta) = \frac{1}{2}(\bar{z}, \bar{\zeta}) + \frac{1}{2}(\hat{z}, \hat{\zeta})$; and (\bar{x}, \bar{y}) , and (\hat{x}, \hat{y}) are obtained by restricting $(\bar{z}, \bar{\zeta})$ and $(\hat{z}, \hat{\zeta})$ to the set of agents \mathcal{M} , it follows that

$$(x, y) = \frac{1}{2}(\bar{x}, \bar{y}) + \frac{1}{2}(\hat{x}, \hat{y}). \quad (4.7)$$

Finally, since $\bar{x}_i^{m^*} = \bar{z}_i^{m^*} = z_i^{m^*} - \epsilon = x_i^{m^*} - \epsilon$ and $\hat{x}_i^{m^*} = \hat{z}_i^{m^*} = z_i^{m^*} + \epsilon = x_i^{m^*} + \epsilon$, it follows that (\bar{x}, \bar{y}) and (\hat{x}, \hat{y}) are feasible solutions that are different than (x, y) . Equation (4.7) implies that (x, y) can be expressed as a convex combination of two feasible solutions. Hence, we conclude that (x, y) cannot be an extreme point of the feasible region of LP2b.

Thus, it follows that if Claim 1 holds, (x, y) such that $x_i^{m^*} \in (0, 1)$ for some m^* and i , cannot be an extreme point of the feasible region of LP2b. We next complete the proof of Lemma 4.3.3, by proving Claim 1. \square

Proof of Claim 1: We will prove the claim via induction over the number of nodes N of the tree. Observe that for $N = 1$, the claim trivially holds by considering solutions \hat{x} and \bar{x} described in the statement of the claim, since when there is only a single node, there are no constraints associated with the edges, and feasibility in LP2b directly follows by observing $\sum_m \hat{x}_i^m = \sum_m \bar{x}_i^m = 1$.

Assume the claim holds for a tree with at most N nodes. We next prove that it holds for $N + 1$ node trees as well. The intuition behind the proof for $N + 1$ node trees can be explained as follows: Assume that for some players m, k , and node i , in the $N + 1$ node tree we have $x_i^m, x_i^k \in (0, 1)$. We consider some node j adjacent to i and modify the initial feasible solution (x, y) by perturbing x_i^m, x_j^m, x_i^k and x_j^k such that the feasibility conditions related to edge (i, j) (i.e., constraints involving variables y_{ij}^l) are satisfied. We consider the tree graphs induced by cutting the original tree graph at this edge. For each of these tree

graphs, the induction hypothesis suggests that it is possible to find solutions that satisfy the claim, and that are consistent with the perturbed values of x_i^m, x_j^m, x_i^k and x_j^k variables. These feasible solutions can be used to construct feasible solutions to the problem for the $N + 1$ node tree, and establish the result. Thus, the main idea behind the proof can be summarized as follows: If the claim holds for tree graphs with at most N nodes, then the claim for any $N + 1$ node tree graph can be established, by adjusting the feasible solution at the end points of an edge, and reducing the problem to constructing new feasible solutions for tree graphs with fewer number of nodes. The induction hypothesis can be used to construct such solutions for tree graphs with fewer number of nodes. Using these, we can construct solutions for the $N + 1$ node tree, satisfying the conditions stated in the statement of the Claim. The proof idea is also illustrated in Figure 4-9.




Figure 4-9: Assume that we can perturb the x_i^m, x_j^m, x_i^k and x_j^k variables in two nodes, i and j , such that feasibility conditions on the edge (i, j) , and its end points are satisfied. Consider the tree graphs induced after deletion of edge (i, j) . For each of the induced subgraphs, a solution that satisfies the feasibility conditions can be constructed using the induction hypothesis. A feasible solution for the $N + 1$ node graph can be constructed using the feasible solutions for these subgraphs.

We next formally prove the claim. Let i be a node chosen as in the statement of the claim. Denote a neighbor of i by j . Consider the subtrees T_1 (containing i), and T_2 (containing j) that can be obtained by deleting edge (i, j) . These trees have fewer than $N + 1$ nodes. Observe that restriction of (x, y) to these subtrees, denoted by $(x, y)_{T_1}$ and $(x, y)_{T_2}$ satisfy the feasibility conditions in LP2b, for these trees.

To prove the claim we make use of Lemma 4.7.1, (proof of this lemma is given after the proof of the claim):

Lemma 4.7.1. *Let $\{x_i^l\}_l, \{x_j^l\}_l$ and $\{y_{ij}^l\}_l$ be a restriction of a solution (x, y) of LP2b to the induced subgraph of nodes $\{i, j\}$ such that*

$$(i) \sum_l x_i^l = 1, \text{ and } \sum_l x_j^l = 1,$$

(ii) *For node i , there are at least two players m, k , such that $x_i^m, x_i^k \in (0, 1)$,*

(iii) $(\{y_{ij}^l\}_l, \{x_i^l\}_l, \{x_j^l\}_l)$ is such that the feasibility conditions associated with edge (i, j) and nodes i and j , are satisfied (i.e., $y_{ij}^l \leq x_i^l, x_j^l$ if $(i, j) \in E^+$ and $0, x_i^l + x_j^l - 1 \leq y_{ij}^l$ if $(i, j) \in E^-, x_i^l, x_j^l \geq 0$).

For some players $b, c \in \mathcal{M}$, constants $\alpha \in [0, 1]$, $z_{ij}^l \in \mathbb{R}$ (for all l), there exists another solution $\{\hat{x}_i^l\}_l, \{\hat{x}_j^l\}$ and $\{\hat{y}_{ij}^l\}_l$ such that for any ϵ where $|\epsilon| > 0$ is sufficiently small, this solution satisfies conditions (i)-(iii), and

1. $\hat{x}_i^m = x_i^m + \epsilon, \hat{x}_i^k = x_i^k - \epsilon, \hat{x}_i^l = x_i^l$ for $l \neq k, m$,
2. $\hat{x}_j^b = x_j^b + \alpha\epsilon, \hat{x}_j^c = x_j^c - \alpha\epsilon$ (where potentially $\{k, m\} \cap \{b, c\} \neq \emptyset$), $\hat{x}_j^l = x_j^l$ for $l \neq b, c$
3. $\hat{y}_{ij}^l = y_{ij}^l + z_{ij}^l\epsilon$.

Consider bidders m, k , node i , and feasible solution (x, y) given in the statement of Claim 1. Observe that for a node j adjacent to i , the feasible solution (x, y) is such that $\{x_i^l\}_l, \{x_j^l\}_l$, and $\{y_{ij}^l\}_l$ satisfy the conditions of Lemma 4.7.1. Thus, using this lemma, for sufficiently small ϵ , we can obtain $\{\hat{x}_i^l\}_l, \{\hat{x}_j^l\}_l$, and $\{\hat{y}_{ij}^l\}_l$, satisfying conditions 1-3 given in the statement of the Lemma. Similarly, replacing ϵ by $-\epsilon$, we can construct $\{\bar{x}_i^l\}_l, \{\bar{x}_j^l\}_l$, and $\{\bar{y}_{ij}^l\}_l$ that satisfy the conditions 1-3 for $-\epsilon$

Focus on the constructed $\{\hat{x}_i^l\}_l$ and $\{\bar{x}_i^l\}_l$. The induction hypothesis suggests that when ϵ is sufficiently small, there exists two feasible solutions associated with the tree T_1 , denoted by $(\hat{x}, \hat{y})_{T_1}$, and $(\bar{x}, \bar{y})_{T_1}$, that are consistent with the constructed $\{\hat{x}_i^l\}_l$ and $\{\bar{x}_i^l\}_l$. In particular, these solutions are such that

- $\hat{x}_i^m = x_i^m + \epsilon, \hat{x}_i^k = x_i^k - \epsilon$
- $\bar{x}_i^m = x_i^m - \epsilon, \bar{x}_i^k = x_i^k + \epsilon$
- $\hat{x}_i^l = \bar{x}_i^l = x_i^l$ for $l \neq m, k$
- $(x, y)_{T_1} = \frac{1}{2}(\bar{x}, \bar{y})_{T_1} + \frac{1}{2}(\hat{x}, \hat{y})_{T_1}$.

Similarly, using the constructed $\{\hat{x}_j^l\}_l$ and $\{\bar{x}_j^l\}_l$, and the induction hypothesis, we conclude that there exists two feasible solutions associated with the tree T_2 , denoted by $(\hat{x}, \hat{y})_{T_2}$, and $(\bar{x}, \bar{y})_{T_2}$ that satisfies (where if $\alpha = 0$, we have $\hat{x} = \bar{x} = x$ and $\hat{y} = \bar{y} = y$)

- $\hat{x}_j^b = x_j^b + \alpha\epsilon, \hat{x}_j^c = x_j^c - \alpha\epsilon$
- $\bar{x}_j^b = x_j^b - \alpha\epsilon, \bar{x}_j^c = x_j^c + \alpha\epsilon$
- $\hat{x}_j^l = \bar{x}_j^l = x_j^l$ for $l \neq b, c$
- $(x, y)_{T_2} = \frac{1}{2}(\bar{x}, \bar{y})_{T_2} + \frac{1}{2}(\hat{x}, \hat{y})_{T_2}$.

Let (\hat{x}, \hat{y}) be a solution to the original problem such that

- (i) it agrees with $(\hat{x}, \hat{y})_{T_1}$ on nodes/edges of T_1 , i.e., for any nodes $n_1, n_2 \in T_1$ and player l , the variables $x_{n_1}^l$ and y_{n_1, n_2}^l have identical values in (\hat{x}, \hat{y}) and $(\hat{x}, \hat{y})_{T_1}$,
- (ii) it agrees with $(\hat{x}, \hat{y})_{T_2}$ on nodes/edges of T_2 ,
- (iii) $\{\hat{y}_{ij}^l\}_l$ is as given above. It is constructed using Lemma 4.7.1, i.e., $\hat{y}_{ij}^l = y_{ij}^l + z_{ij}^l \epsilon$.

Lemma 4.7.1 suggests that (\hat{x}, \hat{y}) is a feasible solution of LP2b (for the original graph). Feasibility of constraints, associated with nodes/edges of T_1 follow since (\hat{x}, \hat{y}) agrees with $(\hat{x}, \hat{y})_{T_1}$, over T_1 . Similarly, feasibility of constraints, associated with nodes/edges of T_2 also follow. Finally, feasibility of the constraints involving edge (i, j) follow, since $\{\hat{y}_{ij}^l\}_l$ are constructed using Lemma 4.7.1, and hence satisfy the feasibility constraints.

Similarly, we denote by (\bar{x}, \bar{y}) a feasible solution to the original problem such that it agrees with $(\bar{x}, \bar{y})_{T_1}$ on nodes of T_1 , it agrees with $(\bar{x}, \bar{y})_{T_2}$ on nodes of T_2 , and $\{\bar{y}_{ij}^l\}_l$ is constructed using Lemma 4.7.1 (i.e., $\bar{y}_{ij}^l = y_{ij}^l - z_{ij}^l \epsilon$).

We claim that (\hat{x}, \hat{y}) and (\bar{x}, \bar{y}) are not only feasible, but also by construction they satisfy the conditions of Claim 1. In particular, $(x, y) = \frac{1}{2}(\bar{x}, \bar{y}) + \frac{1}{2}(\hat{x}, \hat{y})$. This is the case since by construction $(x, y)_{T_1}$ and $(x, y)_{T_2}$ also satisfy a similar condition, and hence for (x, y) this condition immediately holds for all nodes and edges except for (i, j) . For edge (i, j) , on the other hand, this condition follows from Lemma 4.7.1, and the fact that $\hat{y}_{ij}^l = y_{ij}^l + z_{ij}^l \epsilon$ and $\bar{y}_{ij}^l = y_{ij}^l - z_{ij}^l \epsilon$. The first three conditions of the Claim also immediately follow since, by construction $(x, y)_{T_1}$ and $(x, y)_{T_2}$ also satisfy similar conditions. Thus, we establish that for $N + 1$ node graphs, the claim still holds. Therefore, by induction, the claim follows for any tree graph. \square

Proof of Lemma 4.7.1: We prove the lemma by focusing on two main cases: $(i, j) \in E^+$ and $(i, j) \in E^-$. In both cases, we specify α and $\{z_{ij}^l\}$ and use conditions 1-3 of the Lemma to construct the new solutions. Consequently, the solutions we obtain immediately satisfy conditions 1-3, and hence it is sufficient to show that the constructed solutions satisfy conditions (i)-(iii) to complete the proof. Note that the conditions 1 and 2 of the construction and the fact that $\sum_l x_i^l = \sum_l x_j^l = 1$, immediately imply that $\sum_l \hat{x}_i^l = \sum_l \hat{x}_j^l = 1$, i.e., requirement (i) is satisfied by the construction. Similarly since $x_i^m, x_i^k \in (0, 1)$, and $|\hat{x}_i^m - x_i^m|, |\hat{x}_i^k - x_i^k| \leq |\epsilon|$ as a consequence of condition 1 in the construction, it follows that for sufficiently small $|\epsilon|$, $\hat{x}_i^m, \hat{x}_i^k \in (0, 1)$. That is, requirement (ii) is always satisfied. Thus, at each case we consider below, we only need to show that the constructed solution satisfies condition (iii) of the lemma.

Case A: $(i, j) \in E^+$ In this case, we need to choose α, b, c , and $\{z_{ij}^l\}$ so that $\hat{y}_{ij}^l \leq x_i^l, x_j^l$ for all l . We complete the proof by focusing on three subcases:

A1: $x_i^m \neq x_j^m, x_i^k \neq x_j^k$:

Let $\gamma^m = \arg \min_{n \in \{i,j\}} x_n^m$, and $\gamma^k = \arg \min_{n \in \{i,j\}} x_n^k$. In this case we choose $\alpha = 0$ (so choice of b, c does not matter), $z_{ij}^m = 1$ if $\gamma^m = i$, $z_{ij}^k = -1$ if $\gamma^k = i$, and $z_{ij}^l = 0$ otherwise. Note that if $\gamma^m = j$, then $x_j^m < x_i^m$, and hence for sufficiently small ϵ

$$\hat{y}_{ij}^m = y_{ij}^m \leq \hat{x}_j^m = x_j^m \leq \hat{x}_i^m = x_i^m + \epsilon.$$

On the other hand if $\gamma^m = i$, then $x_i^m < x_j^m$, and hence for sufficiently small ϵ

$$\hat{y}_{ij}^m = y_{ij}^m + \epsilon \leq \hat{x}_i^m = x_i^m + \epsilon \leq \hat{x}_j^m = x_j^m.$$

By symmetry it also follows that $\hat{y}_{ij}^k \leq \hat{x}_i^k, \hat{x}_j^k$. These inequalities imply (iii) for $l \in \{m, k\}$. This condition also trivially follows for $l \neq \{m, k\}$, since for these bidders $\hat{x}_i^l = x_i^l, \hat{x}_j^l = x_j^l$, and $\hat{y}_{ij}^l = y_{ij}^l$ by construction.

A2: $x_i^m = x_j^m, x_i^k = x_j^k$:

In this case, we let $\alpha = 1, b = m, c = k$, and $z_{ij}^m = 1, z_{ij}^k = -1$, and $z_{ij}^l = 0$ for $l \neq \{m, k\}$. This construction suggests that

$$\hat{y}_{ij}^m = y_{ij}^m + \epsilon \leq x_i^m + \epsilon = x_j^m + \epsilon \leq \hat{x}_i^m = \hat{x}_j^m. \quad (4.8)$$

By symmetry, the above inequality holds for k by replacing ϵ by $-\epsilon$. These inequalities imply (iii) for $l \in \{m, k\}$. This condition also trivially follows for $l \neq \{m, k\}$, since for these bidders $\hat{x}_i^l = x_i^l, \hat{x}_j^l = x_j^l$, and $\hat{y}_{ij}^l = y_{ij}^l$ by construction.

A3: $x_i^m \neq x_j^m, x_i^k = x_j^k$ or $x_i^m = x_j^m, x_i^k \neq x_j^k$.

Without loss of generality assume that $x_i^m = x_j^m, x_i^k \neq x_j^k$. The other case follows by symmetry.

First assume that $x_i^k < x_j^k$. In this case we choose $\alpha = 1, b = m, c = k$, and $z_{ij}^m = 1, z_{ij}^k = -1$, and $z_{ij}^l = 0$ for $l \neq \{m, k\}$. This construction implies that (4.8) holds for bidder m . On the other hand, for bidder k , we have

$$\hat{y}_{ij}^k = y_{ij}^k - \epsilon \leq x_i^k - \epsilon = \hat{x}_i^k \leq x_j^k - \epsilon = \hat{x}_j^k,$$

where the last inequality follows that ϵ is sufficiently small and $x_i^k < x_j^k$. Thus, we obtain (iii) for $l \in \{m, k\}$. This condition also trivially follows for $l \neq \{m, k\}$, since for these bidders $\hat{x}_i^l = x_i^l, \hat{x}_j^l = x_j^l$, and $\hat{y}_{ij}^l = y_{ij}^l$ by construction.

Next assume that $x_i^k > x_j^k$. In this case, since $\sum_l x_i^l = \sum_l x_j^l = 1$, there exists some bidder r such that $x_i^r < x_j^r$. We choose $\alpha = 1, b = m, c = r$, and $z_{ij}^m = 1$, and $z_{ij}^l = 0$ for $l \neq \{m, k\}$. As before, this construction implies that (4.8) holds for bidder m . On

the other hand, for bidder k we have

$$\hat{y}_{ij}^k = y_{ij}^k \leq \hat{x}_j^k = x_j^k \leq x_i^k - \epsilon = \hat{x}_i^k,$$

where the second inequality holds for sufficiently small ϵ , since $\hat{x}_i^k > \hat{x}_j^k$. By symmetry it follows that $\hat{y}_{ij}^r \leq \hat{x}_i^r, \hat{x}_j^r$ as well. Thus, we obtain (iii) for $l \in \{m, k, r\}$. This condition also trivially follows for $l \notin \{m, k, r\}$, since for these bidders $\hat{x}_i^l = x_i^l, \hat{x}_j^l = x_j^l$, and $\hat{y}_{ij}^l = y_{ij}^l$ by construction.

Case B: $(i, j) \in E^-$ In this case, we need to choose α, b, c , and $\{z_{ij}^l\}$ so that $\hat{x}_i^l + \hat{x}_j^l - 1 \leq \hat{y}_{ij}^l$ and $0 \leq \hat{y}_{ij}^l$ for all l . Once again, we obtain the proof by focusing on three subcases:

B1: $x_i^m + x_j^m \neq 1, x_i^k + x_j^k \neq 1$.

In this case, we choose $\alpha = 0, z_{ij}^m = 1$ if $y_{ij}^m > 0, z_{ij}^k = -1$ if $y_{ij}^k > 0$, and $z_{ij}^l = 0$ otherwise.

Note that this construction implies that for sufficiently small $\epsilon, \hat{y}_{ij}^l \geq 0$. Thus, we only need to establish the constraint $\hat{x}_i^l + \hat{x}_j^l - 1 \leq \hat{y}_{ij}^l$.

Note that if $x_i^m + x_j^m - 1 < 0$, for sufficiently small ϵ , we have $\hat{x}_i^m + \hat{x}_j^m - 1 \leq 0 \leq \hat{y}_{ij}^m$. On the other hand, if $x_i^m + x_j^m - 1 > 0$ (recall that $x_i^m + x_j^m \neq 1$), then it follows from feasibility of (x, y) that $y_{ij}^m > 0$. In this case we have

$$\hat{x}_i^m + \hat{x}_j^m - 1 = x_i^m + x_j^m - 1 + \epsilon \leq y_{ij}^m + \epsilon = \hat{y}_{ij}^m.$$

Thus, for bidder m we have $\hat{x}_i^m + \hat{x}_j^m - 1 \leq \hat{y}_{ij}^m$. Note that by symmetry we also have $\hat{x}_i^k + \hat{x}_j^k - 1 \leq \hat{y}_{ij}^k$, and hence condition (iii) follows for bidders $\{m, k\}$. Finally, since for $l \notin \{m, k\}$ we have $\hat{x}_i^l = x_i^l, \hat{x}_j^l = x_j^l$, and $\hat{y}_{ij}^l = y_{ij}^l$ condition (iii) trivially follows for these bidders as well.

B2: $x_i^m + x_j^m = 1, x_i^k + x_j^k = 1$.

Since $x_i^m, x_i^k \in (0, 1)$, in this case we have $x_j^m, x_j^k \in (0, 1)$. Let $\alpha = 1, b = k, c = m$, and $z_{ij}^l = 0$ for all l . Observe that this construction guarantees that for all $l, \hat{x}_i^l + \hat{x}_j^l = x_i^l + x_j^l$, and $\hat{y}_{ij}^l = y_{ij}^l$. Thus, the conditions of (iii) trivially hold for all l .

B3: $x_i^m + x_j^m = 1, x_i^k + x_j^k \neq 1$, or $x_i^m + x_j^m \neq 1, x_i^k + x_j^k = 1$.

Without loss of generality, we will assume that $x_i^m + x_j^m = 1, x_i^k + x_j^k \neq 1$. The other case follows by symmetry.

Note that $x_i^m + x_j^m = 1, x_i^k + x_j^k \neq 1$ implies that $x_i^k + x_j^k < 1$ since $\sum_l x_i^l = \sum_l x_j^l = 1$. Additionally, this suggests that for some $r \neq m$ (and possibly $r = k$), we have $x_j^r > 0$, and $x_i^r + x_j^r < 1$.

We let $\alpha = 1$, $b = r$, $c = m$, and $z_{ij}^l = 0$ for all l . Note that this construction immediately implies $\hat{y}_{ij}^l = y_{ij}^l \geq 0$ for all l . Thus, we only need to establish the constraint $\hat{x}_i^l + \hat{x}_j^l - 1 \leq \hat{y}_{ij}^l$.

Observe that for $l \neq m$, we have $x_i^l + x_j^l < 1$. Thus, for sufficiently small ϵ , we obtain $\hat{x}_i^l + \hat{x}_j^l - 1 \leq 0 \leq \hat{y}_{ij}^l$. On the other hand, for player m , we have $\hat{x}_i^m + \hat{x}_j^m = x_i^m + x_j^m$ and $\hat{y}_{ij}^m = y_{ij}^m$ by construction. Thus, the constraint follows trivially. Thus, we conclude that the conditions of (iii) hold for all l .

These two cases cover all possible scenarios. Thus, we conclude that the conditions stated in the lemma can always be satisfied for appropriate choice of α, b, c , and $\{z_{ij}^l\}$. \square

Proof of Theorem 4.3.2. Denote the objective value at an optimal solution of LP2 by $OP2$, and the value at an optimal integral solution of LP2 by $OPI2$. Similarly, denote by $OP1$ and $OPI1$ the objective values of an optimal solution of LP1 and an optimal integral solution of LP1. If LP2 has an optimal solution that is integral, we know that $OP2 = OPI2$. Also, since the optimal objective value cannot be larger after imposing the integrality condition, we have $OP1 \geq OPI1$. We next show that $OP1 \leq OP2$ and $OPI2 \leq OPI1$ to establish that $OP1 = OPI1$. Note that this immediately implies that LP1 has an optimal solution that is integral.

$OP1 \leq OP2$: Consider a feasible solution $\{x^m(S)\}$ of LP1. We will show that it is possible to construct a feasible solution of LP2 with the same objective value.

In particular, let $x_i^m = \sum_{S|i \in S} x^m(S)$ for all m and S , and let $y_{ij}^m = \sum_{S|i,j \in S} x^m(S)$. Since the feasible solution of LP1 satisfies $\sum_m \sum_{S|i \in S} x^m(S) \leq 1$, we conclude that

$$\sum_m x_i^m = \sum_m \sum_{S|i \in S} x^m(S) \leq 1.$$

Additionally, since $x^m(S) \geq 0$, we obtain

$$y_{ij}^m = \sum_{S|i,j \in S} x^m(S) \leq \sum_{S|i \in S} x^m(S) \leq x_i^m.$$

Thus, it follows that the constructed solution also satisfies $y_{ij}^m \leq x_i^m, x_j^m$. Finally, for any $ij \in E$ we have

$$\begin{aligned} x_i^m + x_j^m - y_{ij}^m &= \sum_{S|i \in S} x^m(S) + \sum_{S|j \in S} x^m(S) - \sum_{S|i,j \in S} x^m(S) \\ &\leq \sum_{S|i \in S} x^m(S) + \sum_{S|j \in S, i \notin S} x^m(S) \\ &\leq \sum_{S|i \in S} x^m(S) + \sum_{S|i \notin S} x^m(S) \leq 1, \end{aligned}$$

where the last inequality follows since $\sum_S x^m(S) \leq 1$.

Summarizing, we established that the constructed x_i^m, y_{ij}^m is such that it satisfies: (i) $\sum_m x_i^m \leq 1$, (ii) $y_{ij}^m \leq x_i^m, x_j^m$, (iii) $x_i^m + x_j^m - y_{ij}^m \leq 1$. Additionally, since $x^m(S) \geq 0$, we have $x_i^m, y_{ij}^m \geq 0$. Finally, since $\sum_m x_i^m \leq 1$ and $x_i^m \geq 0$, we have $x_i^m \leq 1$, and since $y_{ij}^m \leq x_i^m$ we have $y_{ij}^m \leq 1$. These together imply that $\{x_i^m, y_{ij}^m\}$ is a feasible solution of LP2.

We next show that the optimal objective values of LP1 and LP2 for the solutions $\{x^m(S)\}$ and $\{x_i^m, y_{ij}^m\}$ coincide. Observe that

$$\begin{aligned} \sum_{m,S} x^m(S)v^m(S) &= \sum_{m,S} x^m(S) \left(\sum_{i \in S} w_i^m + \sum_{i,j \in S | ij \in E} w_{ij}^m \right) \\ &= \sum_m \left(\sum_i w_i^m \sum_{S|i \in S} x^m(S) + \sum_{ij \in E} w_{ij}^m \sum_{S|i,j \in S} x^m(S) \right) \quad (4.9) \\ &= \sum_m \left(\sum_i w_i^m x_i^m + \sum_{ij \in E} w_{ij}^m y_{ij}^m \right), \end{aligned}$$

Hence, $\sum_{m,S} x^m(S)v^m(S)$, the objective value of LP1 corresponding to $\{x^m(S)\}$, is equal to $\sum_m \left(\sum_i w_i^m x_i^m + \sum_{ij \in E} w_{ij}^m y_{ij}^m \right)$, the objective value of LP2 corresponding to $\{x_i^m, y_{ij}^m\}$. That is, given a feasible solution of LP1, there exists a corresponding feasible solution of LP2, with the same objective value. Since this is true for the optimal solution of LP1 as well, we conclude $OP1 \leq OP2$.

$OPI2 \leq OPI1$: Consider a feasible integral solution $\{x_i^m, y_{ij}^m\}$ of LP2. Let $S^m = \{i | x_i^m = 1\}$. Since $\sum_m x_i^m \leq 1$, it follows that if $x_i^m = 1$ then $x_i^k = 0$ for all $k \neq m$. Hence, $S^m \cap S^k = \emptyset$ for all $k \neq m$.

Define $\{x^m(S)\}$ such that for all m , $x^m(S^m) = 1$, and $x^m(S) = 0$ for $S \neq S^m$. Observe that such a solution satisfies $\sum_S x^m(S) \leq 1$, and $\sum_m \sum_{S|i \in S} x^m(S) = \sum_m |i \in S^m| x^m(S^m) \leq 1$ (where the latter inequality follows since $S^m \cap S^k = \emptyset$). Thus, it follows that $\{x^m(S)\}$ is a feasible integral solution to LP1.

Note that feasibility of $\{x_i^m, y_{ij}^m\}$ in LP2 implies that if $x_i^m, x_j^m \in \{0, 1\}$, then $y_{ij}^m \in \{0, 1\}$. More precisely for $ij \in E$, if $x_i^m = x_j^m = 1$ then $y_{ij}^m = 1$ (since $x_i^m + x_j^m - 1 \leq y_{ij}^m$). Similarly, if $x_i^m = 0$ then $y_{ij}^m = 0$ (since $y_{ij}^m \leq x_i^m$). This implies that $y_{ij}^m = 1$ if and only if $x_i^m = x_j^m = 1$.

Observe that the construction of $\{x^m(S)\}$ implies that $x_i^m = \sum_{S|i \in S} x^m(S)$. This is because, if $x_i^m = 0$, then $x^m(S) = 0$ for all S containing i , and if $x^m(i)$, there exists exactly one S (denoted by S^m) for which $i \in S^m$ and $x^m(S^m) = 1$. Similarly, our construction implies that $y_{ij}^m = \sum_{S|ij \in S} x^m(S)$. To see this, note that $x_i^m = x_j^m = 1$ if and only if $\sum_{S|i,j \in S} x^m(S) = 1$ (as before if $x_i^m = 0$, then $x^m(S) = 0$ for all $i \in S$, and if $x_i^m = x_j^m = 1$ then there exists exactly one S , denoted by S^m such that $i, j \in S^m$ and $x^m(S^m) = 1$). On

the other hand, it was established before that for $ij \in E$, we have $y_{ij}^m = 1$ if and only if $x_i^m = x_j^m = 1$. These imply that $y_{ij}^m = \sum_{S|ij \in S} x^m(S)$.

Using $x_i^m = \sum_{S|i \in S} x^m(S)$ and $y_{ij}^m = \sum_{S|ij \in S} x^m(S)$, the objective value corresponding to $\{x_i^m, y_{ij}^m\}$ in LP2 (given by $\sum_m \left(\sum_i w_i^m x_i^m + \sum_{ij \in E} w_{ij}^m y_{ij}^m \right)$) and that corresponding to $\{x^m(S)\}$ in LP1 (given by $\sum_{m,S} x^m(S)v^m(S)$) can be related as follows:

$$\begin{aligned}
\sum_m \left(\sum_i w_i^m x_i^m + \sum_{ij \in E} w_{ij}^m y_{ij}^m \right) &= \sum_m \left(\sum_i w_i^m \sum_{S|i \in S} x^m(S) + \sum_{ij \in E} w_{ij}^m \sum_{S|i,j \in S} x^m(S) \right) \\
&= \sum_m \left(\sum_S x^m(S) \sum_{i \in S} w_i^m + \sum_S x_S^m \sum_{ij \in E} w_{ij}^m S \right) \\
&= \sum_{m,S} x^m(S)v^m(S).
\end{aligned} \tag{4.10}$$

Thus, we conclude that given a feasible integral solution of LP2, there exists a corresponding feasible integral solution of LP1, with the same objective value. Since this is true for the optimal integral solution of LP2 as well, we conclude $OPI2 \leq OPI1$.

Summarizing, we have $OP1 \leq OP2$, and $OPI2 \leq OPI1$. Additionally, optimal value is weakly higher without the integrality requirement (i.e., $OPI2 \leq OP2, OPI1 \leq OP1$) and if LP2 has an optimal integral solution, then $OP2 = OPI2$. These imply that

$$OP1 \leq OP2 = OPI2 \leq OPI1 \leq OP1,$$

and hence $OPI1 = OP1$. That is, when LP2 has an optimal solution that is integral, then so does LP1. \square

Chapter 5

General Graphical Valuations

5.1 Introduction and Organization

In the previous chapter, we established the existence of a Walrasian equilibrium, when the underlying value graph is a tree, and the weights bidders associate with the edges of this graph are sign-consistent. We also provided a linear programming formulation that can be used to identify the efficient allocation for such value graphs. Moreover, we established that when we relax the tree graph assumption, or the sign consistency assumption a Walrasian equilibrium need not exist, and the linear programming formulation may have nonintegral solutions. This suggests that if these assumptions do not hold, it may not be possible to design an efficient iterative auction format that relies on anonymous item pricing, and terminates when a “market clearing” condition holds (or a Walrasian equilibrium is reached).

This observation motivates considering more general pricing rules, and iterative auction formats that rely on these pricing rules and terminate when a generalization of the market clearance condition holds. In this chapter, we focus on progressively more general pricing rules, and discuss generalizations of the Walrasian equilibrium (hereafter referred to as pricing equilibria) concept, to settings where such pricing rules are employed. Additionally, we provide linear programming formulations of the efficient allocation problem that have integral optimal solutions, if and only if pricing equilibria with a given pricing rule exist. The LP formulations we obtain in this chapter are stronger than LP1 of Chapter 4. Consequently, one of the LP formulations we provide, has integral optimal solutions (and hence can be used to identify the efficient allocation) for all graphical valuations. This allows for developing efficient iterative auction formats that rely on simple pricing rules, as we discuss in detail in Chapter 6.

More precisely, in this chapter we focus on three new pricing rules that generalize anonymous item pricing of Chapter 4. Recall the definition of anonymous item pricing:

- *Anonymous item pricing* $\{p_i\}_{i \in \mathcal{N}}$: the auctioneer offers a price p_i for every item $i \in \mathcal{N}$, which is the same for all bidders.

In this chapter, we focus on the following generalizations of this pricing rule:

- *Anonymous graphical pricing* $\{p_i, p_{ij}\}_{i \in \mathcal{N}, (i,j) \in E}$: the auctioneer offers a price p_i for every item $i \in \mathcal{N}$, and a discount/markup term (p_{ij}) for every pair of items i, j that are connected by an edge, i.e., $(i, j) \in E$. These prices are the same for all bidders.
- *Bidder-specific item pricing* $\{p_i^m\}_{m \in \mathcal{M}, i \in \mathcal{N}}$: the auctioneer offers a price p_i^m for every item $i \in \mathcal{N}$, and bidder m .
- *Bidder-specific graphical pricing* $\{p_i^m, p_{ij}^m\}_{m \in \mathcal{M}, i \in \mathcal{N}, (i,j) \in E}$: the auctioneer offers a price p_i^m for every item $i \in \mathcal{N}$ and bidder m , and a discount/markup term (p_{ij}^m) for every pair of items i, j that are connected by an edge, and bidder m .

Given these pricing rules, we use $p^m(S^m)$ as a short hand notation for the total price bidder m needs to pay for acquiring bundle S^m of items. For instance for anonymous graphical pricing, this quantity can more explicitly be stated as $p^m(S^m) = \sum_{i \in S^m} p_i + \sum_{ij \in S^m | ij \in E} p_{ij}$, and for other pricing rules it can be defined similarly. We refer to $\{p^m(\cdot)\}$ as the total price function.

We focus on the following generalization of the Walrasian equilibrium (see Definition 4.2.4) concept to these pricing rules.¹

Definition 5.1.1 (Pricing equilibrium). *Let $\{p^m(\cdot)\}$ denote the total price function associated with a pricing rule, and S^m denote the set of items assigned to player m . The tuple $(\{p^m(\cdot)\}_m, \{S^m\}_m)$ is a pricing equilibrium with the given pricing rule if:*

- $p^m(\{i\}) \geq 0$ for every item $i \in \mathcal{N}$, and bidder $m \in \mathcal{M}$,
- S^1, \dots, S^M is a feasible allocation, i.e., $S^k \cap S^m = \emptyset$,
- $v^m(S^m) - p^m(S^m) \geq v^m(S) - p^m(S)$ for every $S \subset \mathcal{N}$,
- $\sum_m p^m(S^m) \geq \sum_m p^m(Z^m)$ for every feasible allocation $\{Z^m\}$.

Observe that for a given pricing rule, the set of items a bidder demands can be stated as $D^m = \arg \max_S v^m(S) - p^m(S)$. This suggests that at a pricing equilibrium the auctioneer can assign every bidder a set of items she demands (hence a market clearance condition holds).

We begin this chapter by studying the pricing equilibrium associated with the anonymous graphical pricing rule in Section 5.2. We provide an LP formulation of the efficient allocation problem, which has an optimal solution that is integral, if and only if a pricing equilibrium with this pricing rule exists. This LP formulation is similar to LP1, and can be used to find the efficient allocation when it has optimal solutions that are integral. On the

¹See (Bikhchandani and Ostroy, 2002) for definition and properties of more general pricing equilibria that potentially rely on offering a different price for every bundle of items.

other hand, it is stronger, in the sense that it has constraints associated not only with the nodes of the underlying value graph, but also with its edges. We show that when the underlying value graph has a 5-clique (as a minor), an optimal solution of this LP formulation may give the efficient outcome, while optimal solutions of LP1 fail to do so. Conversely, we establish that if the underlying graph does not have a 4-clique (as a minor), an optimal solution of this LP formulation gives the efficient outcome if and only if an optimal solution of the LP formulation associated with anonymous item pricing does so. These results suggest that this LP formulation cannot be used to identify the efficient outcome, even for simple examples provided in Section 4.4, and a pricing equilibrium with this pricing rule may not exist. Thus, it is necessary to focus on more general pricing rules and pricing equilibria, for the design of iterative auctions that guarantee efficiency for general graphical valuations.

This motivates us to consider the bidder-specific generalizations of the anonymous item/graphical pricing rules. In Section 5.3, we focus on the pricing equilibrium associated with the bidder-specific item pricing rule, and provide an LP formulation that has integral optimal solutions when pricing equilibrium with this pricing rule exists. We establish that this LP formulation can find the efficient outcome if and only if LP1 can. That is, considering this LP (and the associated more complex pricing rule) does not allow us to identify the efficient outcome for a larger class of graphical valuations. In contrast, in Section 5.4, we establish that a pricing equilibrium associated with the bidder-specific graphical pricing rule always exists. We also provide an LP formulation associated with this pricing rule, and show that this formulation has integral optimal solutions (and identifies the efficient outcome) for all graphical valuations.

Importantly, in Section 5.5, we establish that the structure of the LP formulation introduced in Section 5.4 can be systematically generalized to provide LP formulations that obtain the efficient outcome for a generalization of graphical valuations, namely additively decomposable valuations. Additively decomposable valuations allow for complementarity/substitutability not only between pairs but also among larger subsets of items as well. The LP formulation (and the associated pricing rule) we provide in this section becomes progressively stronger as the underlying additively decomposable valuation functions become more general.

The results of this chapter provide a way of systematically obtaining LP formulations that can find the efficient allocation for all graphical valuations (and additively decomposable valuations). Additionally, these results suggest that efficient iterative auction formats that rely on bidder-specific graphical pricing, and terminate at a pricing equilibrium can be used to implement the efficient outcome for general graphical valuations. We provide such an iterative auction format in Chapter 6.

We close this chapter with a summary of its main contributions in Section 5.6. Some of the proofs are delegated to Section 5.7, in order to simplify the exposition.

5.2 Anonymous Graphical Pricing

In this section, we focus on pricing equilibria with the anonymous graphical pricing rule, and characterize its properties. We start this section by providing an LP formulation of the efficient allocation problem, whose dual suggests this pricing rule. We state our formulation by making use of the notion of complete feasible allocation introduced in Definition 4.2.3. Recall that the set of complete feasible allocations is denoted by χ . Using this notation, the LP formulation that we focus on in this section is stated below:

$$\begin{aligned}
 & \max \sum_m \sum_S x^m(S) v^m(S) \\
 & \text{s.t.} \sum_S x^m(S) \leq 1 \quad \forall m \\
 LP3 : & \sum_m \sum_{S|i \in S} x^m(S) \leq \sum_m \sum_{\mu \in \chi | i \in \mu^m} \delta_\mu \quad \forall i \\
 & \sum_m \sum_{S|ij \in S} x^m(S) = \sum_m \sum_{\mu \in \chi | ij \in \mu^m} \delta_\mu \quad \forall ij \in E \\
 & \sum_{\mu \in \chi} \delta_\mu = 1 \\
 & \delta_\mu \geq 0, x^m(S) \geq 0.
 \end{aligned}$$

In this LP formulation, similarly to LP1 of Chapter 4, the variable $x^m(S)$ captures an assignment of bundle S to bidder m . On the other hand, unlike LP1, in LP3 we have a variable δ_μ for every complete allocation $\mu \in \chi$. The variable δ_μ captures the allocation that is preferred by the auctioneer, i.e., $\delta_\mu = 1$ if the auctioneer would like to assign bundle μ^m to bidder m , for every $m \in \mathcal{M}$.

Intuitively, the first constraint in LP3 suggests that each bidder m receives at most one bundle S . The second constraint suggests that an item i is assigned to some bidder if it belongs to a complete allocation that is preferred by the auctioneer. Observe that by definition, for any complete allocation μ , every item i is such that $i \in \mu^m$ for some bidder m . Since $\sum_{\mu \in \chi} \delta_\mu = 1$, this implies that $\sum_m \sum_{\mu \in \chi | i \in \mu^m} \delta_\mu = 1$. That is, the second constraint is equivalent to $\sum_m \sum_{S|i \in S} x^m(S) \leq 1$, which is also present in LP1. The third constraint implies that if the auctioneer prefers to assign a pair of items (ij) together, then the assignment of items $(\{x^m(S)\})$ should be consistent with this preference.

Observe that the objective of this LP formulation is exactly the same as in LP1, i.e., maximizing efficiency. LP3 is closely related to LP1 but it is stronger, in the sense that the projection of the set of feasible solutions of LP3 to the space of variables $\{x^m(S)\}$ is a subset of the feasible set of LP1. On the other hand, any complete feasible allocation still corresponds to a feasible solution of LP3, as we formalize in the next Lemma.

Lemma 5.2.1. *(i) Let $(\{x^m(S)\}_{m,S}, \{\delta_\mu\}_\mu)$ be a feasible solution of LP3, then $\{x^m(S)\}$ is a feasible solution of LP1.*

(ii) Let $\hat{\mu}$ be a complete feasible allocation, then $x^m(S) = 0$ if $S \neq \hat{\mu}^m$, $x^m(\hat{\mu}^m) = 1$ and $\delta_\mu = 0$ if $\mu \neq \hat{\mu}$, $\delta_{\hat{\mu}} = 1$ is a feasible solution of LP3.

Proof. The feasibility of $(\{x^m(S)\}_{m,S}, \{\delta_\mu\}_{\mu \in \chi})$ in LP3 implies that $\sum_m \sum_{\mu \in \chi | i \in \mu^m} \delta_\mu \leq 1$. Consequently, it follows that $\sum_S x^m(S) \leq 1$ for all m , and $\sum_m \sum_{S | i \in S} x^m(S) \leq 1$ for all i . Together with nonnegativity constraints, this implies that $\{x^m(S)\}_{m,S}$ is feasible in LP1, and the first part of the lemma follows. The second part can be immediately verified from LP3. \square

Note that under Assumption 4.2.1, there exists an efficient allocation that is also complete. Together with the above lemma, this observation suggests that if LP3 has an optimal solution that is integral, then in this solution $\delta_{\hat{\mu}} = 1$ for some $\hat{\mu}$, and the corresponding allocation $\hat{\mu}$ is efficient. These imply that for classes of valuation functions for which LP3 has optimal solutions that are integral, this LP formulation can be used to identify the efficient outcome. Moreover, the first part of the lemma implies that the optimal objective value of LP3 is weakly lower than that of LP1. Hence, there may be classes of valuation functions, where LP3 can be used to identify the efficient allocation, while LP1 cannot.

We next provide the dual of LP3:

$$\begin{aligned}
 & \min \pi^s + \sum_m \pi^m \\
 & \text{s.t. } \pi^m \geq v^m(S) - \sum_{i \in S} p_i - \sum_{i,j \in S | ij \in E} p_{ij} \quad \forall S, m \\
 D3 : & \pi^s \geq \sum_m \left(\sum_{i \in \mu^m} p_i + \sum_{i \in \mu^m} p_{ij} \right) \quad \forall \mu \in \chi \\
 & p_i, \pi^m \geq 0.
 \end{aligned}$$

Observe that in the dual problem we have a price variable for each node (item) i , and edge (pair of items) $ij \in E$. The quantity $\sum_{i \in S} p_i + \sum_{i,j \in S | ij \in E} p_{ij}$ can be interpreted as the price a bidder needs to pay for acquiring a bundle S of items at the given node and edge prices. Intuitively, this corresponds to offering prices for each item, and complementing these prices with additional discounts/markups for pairs of items. The variable π^m is an upper bound on the surplus $(v^m(S) - \sum_{i \in S} p_i - \sum_{i,j \in S | ij \in E} p_{ij})$ a bidder can raise by acquiring the bundle S of items at the given prices. Thus, in a similar fashion to D1, it can be interpreted as the surplus of bidder m . The variable π^s , on the other hand, is an upper bound on the quantity $\sum_m \left(\sum_{i \in \mu^m} p_i + \sum_{i \in \mu^m} p_{ij} \right)$ for all complete allocations μ . The latter quantity is the revenue the auctioneer can raise by assigning items according to allocation μ at the given prices. At optimality π^s is equal to the maximum revenue the auctioneer can raise at those prices. This implies that this variable can be interpreted as the revenue of the auctioneer.

By Assumption 4.2.1, it follows that LP3 has an optimal solution that is integral if and only if it has such a solution that also corresponds to a complete feasible allocation. Note that LP3 has such an optimal solution if and only if there exists some allocation $\{S^m\}$ such that (i) $x^m(S^m) = 1$ and $x^m(S) = 0$ for $S \neq S^m$, (ii) $\delta_\mu = 1$ for $\mu = \{S^m\}$ and $\delta_\mu = 0$ for $\mu \neq \{S^m\}$. However by complementary slackness this is equivalent to the existence of a dual feasible solution such that $\pi^m = v^m(S^m) - p^m(S^m)$ for all m , and $\pi^s = \sum_m \sum_{i \in S^m} p^m(S^m)$. On the other hand, this immediately implies that LP3 has an integral optimal solution if and only if a pricing equilibrium with anonymous graphical pricing rule exists. This suggests that by using iterative algorithms for the solutions of LP3 and D3, it is possible to develop new iterative auctions that rely on anonymous graphical pricing rule, and terminate when a pricing equilibrium with this pricing rule is identified.

We next characterize the class of valuation functions for which LP3 has optimal solutions that are integral, while LP1 does not. Before we state our result we introduce some necessary definitions.

Definition 5.2.1 (Forest). *An undirected graph is a forest if its connected components are trees.*

Definition 5.2.2 (Series-Parallel Graph (Chopra (1994))). *A graph is a series-parallel graph if it can be obtained from a forest by repeatedly adding an edge in parallel to an existing one or by replacing an edge by a path.*

Observe that this definition allows the graph to have multiple edges between two nodes. Such graphs are sometimes referred to as *multigraphs*. Graphs that are not multigraphs are referred to as *simple graphs*. In this thesis we restrict attention only to simple graphs (even when we focus on series-parallel graphs).

Definition 5.2.3 (Minor). *An undirected graph G_m is called a minor of the graph G if G_m can be obtained from G by deletion of edges, vertices, and contraction of edges.*

Definition 5.2.4 (Complete Graph). *A graph is a complete graph if there is an edge between any two pair of vertices. A complete graph with n nodes is denoted by K_n .*

In our exposition K_4 and K_5 play a key role. It is known that these graphs are closely related to series-parallel graphs. In particular, series-parallel graphs are equivalent to graphs that do not contain K_4 as a minor (Chopra, 1994; Diestel, 2005; Duffin, 1965).

We start our analysis of LP3 by providing an example where LP3 has optimal solutions that are integral, but LP1 does not. That is, for this example the efficient allocation can be found by using the stronger formulation in LP3, but not the one in LP1.

Example 5.2.1. *Let the underlying value graph be K_5 , and assume that there are 3 bidders. We assume that all players have the same set of weights, and these weights are as given in Figure 5-1.*




Figure 5-1: All players have the same weights, and the edge weights are as given in the above figure. In order to simplify the figure, some weights are omitted. In particular, the dashed lines have weight -30 , and all nodes have weight 100 .

In this example, the efficient allocation is obtained by assigning item A to the first bidder, items B and E to the second one, and items C and D to the third one. The total welfare corresponding to this allocation is 525 . On the other hand, $LP1$ has a feasible solution with higher objective value. Namely: $x^1(\{A, B\}) = x^1(\{A\}) = 1/2$, $x^2(\{B, E, C\}) = x^2(\emptyset) = 1/2$, $x^3(\{E, D\}) = x^3(\{D, C\}) = 1/2$ and the remaining $x^m(S)$ variables are equal to zero. It can be checked that the objective value of $LP1$ associated with this solution is 527 . Thus, $LP1$ does not have an integral optimal solution for this example.

On the other hand, a feasible solution to $D3$ is obtained by setting $p_i = w_i^m$, and $p_{ij} = w_{ij}^m$ for all nodes i and edges ij , $\pi^m = 0$ for all m . It can be immediately checked from $D3$ that by setting π^s equal to the maximum welfare that can be associated with this example (525), the feasibility of the constructed dual solution follows. Hence, we have a dual feasible solution with objective value 525 . This suggests that the maximum objective value of $LP3$ is bounded by the value of the dual feasible solution, 525 . On the other hand, Lemma 5.2.1 suggests that the efficient allocation corresponds to a feasible solution of $LP3$ with this objective value. Thus, it follows that $LP3$ has an optimal solution that is integral.

Note that in this example we chose the node weights large, in order to assure that monotonicity of valuation functions (Assumption 4.2.1) holds. However, it is possible to choose smaller node weights to construct examples where the gap between the optimal objective values of $LP1$ and $LP3$ is larger.

This example shows that by employing $LP3$ it may be possible to find the efficient outcome in cases where using $LP1$ does not suffice. Our next result provides general conditions, under which this might be the case.

Theorem 5.2.1. (i) Assume that there are at least three bidders. For value graphs that are series-parallel (or that exclude a K_4 as a minor) $LP3$ has an optimal solution that is integral if and only if $LP1$ does.

(ii) For value graphs that have a K_5 minor, for some choice of weights, LP3 has an optimal solution that is integral, whereas LP1 does not.

The proof of this theorem is provided in Section 5.7. Since integral feasible solutions of LP1 and LP3 correspond to feasible allocations, this result implies that, for graphs that include K_5 as a minor, using LP3 may allow for finding the efficient outcome, even when using LP1 does not. On the other hand, it also suggests that for graphs that have a simple structure, such as series-parallel graphs (or graphs that do not include K_4 as a minor), using the more complex LP formulation (LP3) is not valuable, in the sense that this does not allow us to find the efficient outcome for a larger class of graphical valuation functions. Importantly, this implies that LP3 is not sufficient to find the efficient outcome, even for the simple examples provided in Section 4.4, as the value graphs in these examples do not have K_4 as a minor, and in these examples LP1 cannot find the efficient outcome. Since a pricing equilibrium with anonymous graphical pricing exists only when LP3 has optimal solutions that are integral, it follows that for the aforementioned examples a pricing equilibrium with this pricing rule cannot exist. Thus, in order to be able to develop efficient iterative auctions that terminate when a pricing equilibrium is reached, we need to consider more general pricing rules. We explore such pricing rules in the subsequent sections.

5.3 Bidder-Specific Item Pricing

In this section, we focus on the bidder-specific item pricing rule, and provide an LP formulation of the efficient allocation problem that has integral optimal solutions if and only if a pricing equilibrium with this pricing rule exists. The main result of this section suggests that the LP formulation associated with bidder-specific item prices can find the efficient outcome if and only if the formulation associated with anonymous item prices can. This suggests that a more general pricing rule may be necessary for the existence of a pricing equilibrium, and LP formulations that identify the efficient allocation for general graphical valuations.

The LP formulation that we focus on in this section (LP4), and its dual (D4) are presented below:

$$\begin{aligned}
 & \max \sum_m \sum_S x^m(S) v^m(S) \\
 & \text{s.t.} \sum_S x^m(S) \leq 1 \quad \forall m \\
 \text{LP4 :} \quad & \sum_{S|i \in S} x^m(S) \leq \sum_{\mu \in \chi | i \in \mu^m} \delta_\mu \quad \forall i, m \\
 & \sum_{\mu \in \chi} \delta_\mu = 1 \\
 & \delta_\mu \geq 0, x^m(S) \geq 0.
 \end{aligned}$$

$$\begin{aligned}
& \min \pi^s + \sum_m \pi^m \\
& \text{s.t. } \pi^m \geq v^m(S) - \sum_{i \in S} p_i^m \quad \forall S, m \\
\text{D4: } & \pi^s \geq \sum_m \left(\sum_{i \in \mu^m} p_i^m \right) \quad \forall \mu \in \chi \\
& p_i^m, \pi^m \geq 0.
\end{aligned}$$

Observe that LP4 is similar to LP3 in that it relies on $\{\delta_\mu\}$ variables that capture the allocation preferred by the auctioneer. On the other hand, it imposes the second constraint for all players m separately (as opposed to imposing a single constraint after summation over all m , as LP3 does). Additionally, it does not impose constraints for the edges of the underlying graph. Consequently, D4 has bidder-specific price (p_i^m), and bidder surplus (π^m) variables, and a seller revenue variable π^s . However, it does not associate any dual variables with the edges of the underlying value graph.

By Assumption 4.2.1, it follows that if LP4 has an optimal solution that is integral, it also has such a solution that corresponds to a complete allocation. It can be seen that LP4 has such a solution if and only if there exists some allocation $\{S^m\}$ such that (i) $x^m(S^m) = 1$ and $x^m(S) = 0$ for $S \neq S^m$, (ii) $\delta_\mu = 1$ for $\mu = \{S^m\}$ and $\delta_\mu = 0$ for $\mu \neq \{S^m\}$. However by complementary slackness this is equivalent to existence of a dual feasible solution to D4 such that $\pi^m = v^m(S^m) - p^m(S^m)$, and $\pi^s = \sum_m \sum_{i \in S^m} p_i^m(S^m)$. On the other hand, this immediately implies that LP4 has an optimal solution that is integral if and only if a pricing equilibrium with bidder-specific item pricing rule exists.

We next establish that LP4 is also closely related to LP1. In particular, we show that LP4 has optimal solutions that are integral if and only if LP1 does.

Theorem 5.3.1. *LP4 has optimal solutions that are integral if and only if LP1 does.*

Proof. Consider any feasible solution $\{x^m(S), \delta_\mu\}$ of LP4. Observe that

$$\sum_m \sum_{S|i \in S} x^m(S) \leq \sum_m \sum_{\mu \in \chi | i \in \mu^m} \delta_\mu \leq \sum_{\mu \in \chi} \delta_\mu \leq 1.$$

Thus, it follows that $\{x^m(S)\}$ is feasible in LP1. Moreover, since LP1 and LP4 have the same objective function, this implies that the optimal objective value of LP4 (denoted by OPT4) is weakly lower than that of LP1 (denoted by OPT1),

$$\text{OPT4} \leq \text{OPT1}. \tag{5.1}$$

Consider the dual optimal solution $\{\pi^m, \pi^s, p_i^m\}$ of D4. Observe that if $p_i^m \geq p_i^k$, then for any $\mu \in \chi$ such that $i \in \mu^k$, we have $\sum_l \sum_{i \in \mu^l} p_i^l \leq \sum_l \sum_{i \in \hat{\mu}^l} p_i^l$, where $\hat{\mu} \in \chi$ is such that $\hat{\mu}^k = \mu^k - \{i\}$, $\hat{\mu}^m = \mu^m \cup \{i\}$, and $\hat{\mu}^l = \mu^l$ for $l \neq k, m$. This observation suggests that the

constraint involving π^s is active for μ such that for all items $i \in \mu^m$ and bidders k we have $p_i^m \geq p_i^k$. Thus, we obtain $\pi^s = \sum_i \max_k p_i^k$. We construct another dual feasible solution $(\bar{\pi}^m, \bar{\pi}^s, \bar{p}_i^m)$ by setting $\bar{p}_i^m = \max_k p_i^k$, $\bar{\pi}^s = \pi^s$, and $\bar{\pi}^m = \pi^m$. The solution satisfies the first constraint of D4, since in our construction, we weakly increase p_i^m , while keeping π^m intact. The feasibility of the second constraint follows since, by dual optimality of the original solution we have $\pi^s = \sum_i \max_k p_i^k$. Observe that since $\bar{\pi}^s = \pi^s$, and $\bar{\pi}^m = \pi^m$ the new solution has the same objective value as the original one, and hence is also optimal in D4. In addition, the solution $\hat{\pi}^m = \bar{\pi}^m$, $\hat{p}_i = \bar{p}_i^m$ is a feasible solution of D1 with the same objective value. Thus, given an optimal solution of D4, we can construct a feasible solution of D1 with the same objective value. This suggests that the optimal objective value of LP1 is weakly lower than that of LP4.

$$OPT1 \leq OPT4. \tag{5.2}$$

It follows from (5.1) and (5.2) that LP1 and LP4 have the same optimal objective value. Assume that one of them has an integral solution that is optimal. Note that there exists an allocation associated with this solution (obtained in particular by assigning bundle S to bidder m if and only if $x^m(S) = 1$), and the optimal objective value of this problem is equal to the welfare associated with this allocation. Moreover, Assumption 4.2.1 implies that there also exists a complete feasible allocation $\{\hat{S}^m\}$ that leads to the same welfare. The solution $x^m(S) = 1$ for $S = \hat{S}^m$ and $x^m(S) = 0$ otherwise is a feasible solution of LP1 that achieves the optimal welfare. Similarly, complementing this solution with $\delta_\mu = 0$ if $\mu \neq \{\hat{S}^m\}$ and $\delta_\mu = 1$ otherwise we obtain a feasible solution of LP4 achieving the same welfare. Thus, we conclude that if one problem has an optimal solution that is integral, then we can construct a feasible solution to the other that achieves the same objective value. By (5.1) and (5.2) we conclude that this is an optimal solution of the latter problem, and the claim follows.

□

Note that nonintegral solutions of LP4 do not correspond to feasible allocations of items. Thus, our result implies that LP4 cannot be used to find the efficient allocation in cases where LP1 fails to do so. Additionally, a pricing equilibrium with this pricing rule exists, only when a Walrasian equilibrium exists (and hence LP1 has optimal solutions that are integral). Therefore, in order to be able to find the efficient outcome for general classes of graphical valuations, we need more general pricing rules, and LP formulations. We provide one such formulation in the next section.

5.4 Bidder-Specific Graphical Pricing

In this section, we focus on bidder-specific graphical pricing, and provide an LP formulation that has integral optimal solutions if and only if a pricing equilibrium with this pricing rule exists. This LP formulation is stronger than all of the LP formulations we discussed so far in this chapter. We establish that for all graphical valuations, this LP formulation has an optimal solution that is integral, and hence can be used to identify the efficient outcome. Our result also implies that for graphical valuations a pricing equilibrium with this pricing rule always exists.

We start this section by providing the LP formulation (LP5) that we focus on in the rest of this section, and its dual (D5).

$$\begin{aligned}
 & \max \sum_m \sum_S x^m(S) v^m(S) \\
 & \text{s.t.} \sum_S x^m(S) \leq 1 \quad \forall m \\
 LP5 : & \sum_{S|i \in S} x^m(S) \leq \sum_{\mu \in \chi | i \in \mu^m} \delta_\mu \quad \forall i, m \\
 & \sum_{S|ij \in S} x^m(S) = \sum_{\mu \in \chi | ij \in \mu^m} \delta_\mu \quad \forall ij \in E, m \\
 & \sum_{\mu \in \chi} \delta_\mu = 1 \\
 & \delta_\mu \geq 0, x^m(S) \geq 0.
 \end{aligned}$$

$$\begin{aligned}
 & \min \pi^s + \sum_m \pi^m \\
 & \text{s.t.} \pi^m \geq v^m(S) - \sum_{i \in S} p_i^m - \sum_{i,j \in S | ij \in E} p_{ij}^m \quad \forall S, m \\
 D5 : & \pi^s \geq \sum_m \left(\sum_{i \in \mu^m} p_i^m + \sum_{i \in \mu^m} p_{ij}^m \right) \quad \forall \mu \in \chi \\
 & p_i^m, \pi^m \geq 0.
 \end{aligned}$$

Observe that the main difference between LP5 and LP3, is that the node and edge constraints (the second and third constraints) are imposed for every bidder m separately in LP5, whereas they are aggregated over different bidders in LP3. This implies that every feasible solution of LP5 is a feasible solution of LP3, i.e., LP5 is a stronger formulation. The presence of different node/edge constraint for each bidder, on the other hand, leads to bidder-specific price variables p_i^m and p_{ij}^m in the dual problem.

Similar to LP3, LP5 is closely related to the existence of a pricing equilibrium. As before, if LP5 has an optimal solution that is integral, then (by Assumption 4.2.1) it has an optimal solution that corresponds to a complete allocation. That is, there exists some allocation

$\{S^m\}$ corresponding to an optimal solution such that (i) $x^m(S^m) = 1$ and $x^m(S) = 0$ for $S \neq S^m$, (ii) $\delta_\mu = 1$ for $\mu = \{S^m\}$ and $\delta_\mu = 0$ for $\mu \neq \{S^m\}$. Moreover, due to complementary slackness, in this setting a dual feasible solution such that $\pi^m = v^m(S^m) - p^m(S^m)$, and $\pi^s = \sum_m \sum_{i \in S^m} p^m(S^m)$ always exists. Note that this is equivalent to existence of a pricing equilibrium. Hence, we conclude that the existence of a pricing equilibrium with bidder-specific graphical pricing is equivalent to the presence of integral optimal solutions in LP5.

We next establish that provided that the valuations are graphical, this stronger formulation has optimal solutions that are integral.

Theorem 5.4.1. *Assume that bidders have graphical valuations. Then, LP5 has an optimal solution that is integral.*

Proof. Assume that a complete feasible allocation $\hat{\mu}$ is given. Consider the corresponding feasible integral solution of LP5 that is obtained by choosing $\delta_{\hat{\mu}} = 1$, and $x^m(\hat{\mu}^m) = 1$ for all m , and setting $\delta_\mu = x^m(S) = 0$ for $\mu \neq \hat{\mu}$, $S \neq \hat{\mu}^m$. Observe that the objective value of LP5 at this solution is equal to the welfare associated with feasible allocation $\hat{\mu}$. Since this observation holds for any complete feasible allocation, it follows that the optimal objective value of LP5 is at least the maximum welfare obtained by an allocation. Denote the maximum welfare by W^* .

Consider the following dual solution: $p_i^m = w_i^m$, $p_{ij}^m = w_{ij}^m$, $\pi^m = 0$, and $\pi^s = W^*$. By construction this solution suggests $v^m(S) - \sum_{i \in S} p_i^m - \sum_{i,j \in S | ij \in E} p_{ij}^m = 0$, thus the first constraint of D5 is immediately satisfied. Additionally, the construction also implies that $W^* \geq \sum_m \left(\sum_{i \in \mu^m} p_i^m + \sum_{i,j \in \mu^m} p_{ij}^m \right)$ for all complete feasible allocations μ . Thus, we conclude that $\pi^s = W^*$ satisfies the second constraint, and feasibility of the suggested solution follows.

Observe that the objective value associated with the feasible dual solution we constructed is W^* . However, as we established earlier, the optimal objective value of LP5 is at least W^* . Thus, it follows that the dual feasible solution we constructed is optimal, and the feasible integral solution of LP5 that corresponds to the welfare maximizing allocation is optimal in the primal problem. \square

Note that any complete feasible allocation $\hat{\mu}$ corresponds to a feasible integral solution of LP5 (where $x^m(\hat{\mu}^m) = 1$ for all m , $\delta_{\hat{\mu}} = 1$, and $x^m(S) = \delta_\mu = 0$ for remaining variables). Thus, Theorem 5.4.1 implies that the efficient allocation can always be found by solving LP5. Moreover, for graphical valuations a pricing equilibrium with bidder-specific graphical pricing always exists. Interestingly, this result relies on having a dual problem which suggests a pricing rule that has the same structure as the graphical valuations themselves, i.e., a pricing rule that involves a bidder-specific price variable for each node/edge of the underlying value graph. As we discuss in Section 5.5, this observation is more general, and often it is the case that the efficient outcome can be found by solving a linear program, whose dual suggests the same pricing “structure” as the valuation functions.

5.5 Additively Decomposable Valuations

In this section, we extend the results of Section 5.4 to additively decomposable valuations, which is a class of valuation functions that generalize graphical valuations. The valuations in this class allow for complementarity/substitutability not only for pairs, but also for larger sets of items. Moreover, they exhibit an additive structure over subsets of items. We establish that in order to find the efficient outcome for additively decomposable valuation functions, it is sufficient to solve an LP formulation whose dual suggests a pricing rule with the same additively decomposable structure. Our result allows for systematically formulating the efficient allocation problem as a linear program that has a simple structure, whenever the underlying valuation functions exhibit a simple structure.

We start our analysis by formally defining additively decomposable valuations. Consider a collection of subset of items \mathcal{B} , i.e., $B \in \mathcal{B}$ is such that $B \subset \mathcal{N}$. Assume that the valuations of bidders can be additively decomposed over these subsets as follows:

$$v^m(S) = \sum_{B \in \mathcal{B}} w_B^m(S \cap B), \quad (5.3)$$

where $w_B^m : 2^B \rightarrow \mathbb{R}$, captures the component of the valuation of bidder m associated with subset B . We refer to such valuations as *additively decomposable valuations* with collection \mathcal{B} .

We note that any valuation function can be represented using additively decomposable valuations, by considering a collection \mathcal{B} such that $\mathcal{N} \in \mathcal{B}$. On the other hand, if \mathcal{B} consists of few sets of small cardinality, then the valuation functions can be compactly represented by specifying their components $\{w_B^m\}$. For instance, graphical valuations are a special class of additively decomposable valuations, where

- \mathcal{B} consists of singletons, and pairs of items that correspond to the edges of the underlying value graph,
- $w_i^m(S) = 0$ if $S \neq \{i\}$, and it equals to the weight associated with node i otherwise,
- $w_{ij}^m(S) = 0$ if $S \neq \{i, j\}$, and it equals to the edge weight for edge (i, j) otherwise.

Observe that in graphical valuations node weights are nonnegative, whereas edge weights can be negative or positive. In order to capture a similar structure in additively decomposable valuations, we let $\mathcal{B} = \mathcal{B}_+ \cup \mathcal{B}_0 \cup \mathcal{B}_-$, where \mathcal{B}_+ , \mathcal{B}_0 , \mathcal{B}_- are mutually exclusive sets such that $w_B^m(\cdot) \geq 0$ for $B \in \mathcal{B}_+$, $w_B^m(\cdot) \leq 0$ for $B \in \mathcal{B}_-$ and $w_B^m(\cdot)$ can be positive or negative for $B \in \mathcal{B}_0$.

We next provide a generalization of LP5 (denoted by *LP5G*) that allows for finding the

efficient allocation for additively decomposable valuations, as we establish subsequently.

$$\begin{aligned}
& \max \quad \sum_m \sum_S x^m(S) v^m(S) \\
& \text{s.t.} \quad \sum_S x^m(S) \leq 1 \quad \forall m \\
& \sum_{S|S'=S \cap B} x^m(S) = \sum_{\mu \in \chi | \mu^m \cap B = S'} \delta_\mu \quad \forall m, B \in \mathcal{B}_0, S' \subset B \\
LP5G : \quad & \sum_{S|S'=S \cap B} x^m(S) \leq \sum_{\mu \in \chi | \mu^m \cap B = S'} \delta_\mu \quad \forall m, B \in \mathcal{B}_+, S' \subset B \\
& \sum_{S|S'=S \cap B} x^m(S) \geq \sum_{\mu \in \chi | \mu^m \cap B = S'} \delta_\mu \quad \forall m, B \in \mathcal{B}_-, S' \subset B \\
& \sum_{\mu \in \chi} \delta_\mu \leq 1 \\
& x^m(S), \delta_\mu \geq 0 \quad \forall m, S, \mu.
\end{aligned}$$

The corresponding dual LP (D5G) is given as follows:

$$\begin{aligned}
& \min \quad \pi^s + \sum_m \pi^m \\
& \text{s.t.} \quad \pi^s \geq \sum_m \sum_B p_B^m(S^m \cap B) \quad \forall \{S^m\} \\
D5G : \quad & \pi^m \geq v^m(S) - \sum_B p_B^m(S \cap B) \quad \forall m, S \\
& \pi^m \geq 0 \quad \forall m \\
& p_B^m(\cdot) \geq 0 \text{ for } B \in B_+, p_B^m(\cdot) \leq 0 \text{ for } B \in B_-.
\end{aligned}$$

As mentioned earlier, for graphical valuations, we can focus on a collection \mathcal{B} that consists only of singletons, and pairs of items that are connected with an edge in the underlying value graph. Moreover, the node weights are nonnegative, whereas edge weights can be negative or positive. Thus, the components of the valuation function that are associated with node weights belong to \mathcal{B}_+ , whereas the components associated with edge weights belong to \mathcal{B}_0 . These observations suggest that LP5G reduces to LP5 for graphical valuations.

Our next theorem establishes that the primal LP always has integral optimal solutions, and can be used to identify the efficient outcome for additively decomposable valuations.

Theorem 5.5.1. *Assume that bidders have additively decomposable valuations. LP5G has an optimal solution that is integral.*

Proof. The proof is analogous to that of Theorem 5.4.1. It immediately follows by (i) establishing the feasibility of the dual solution $p_B^m = w_B^m$, $\pi^m = 0$, $\pi^s = W^*$ (where W^* denotes the maximum welfare that can be associated with a feasible allocation), and (ii)

showing that there exists an integral primal feasible solution (that coincides with the efficient allocation) which also achieves W^* . Note that in establishing dual feasibility, the sign constraints do not play a role, since by construction $p_B^m = w_B^m \leq 0$ for $B \in \mathcal{B}_-$, and $p_B^m = w_B^m \geq 0$ for $B \in \mathcal{B}_+$. \square

We conclude this section by noting an interesting implication of the above theorem: If the valuations exhibit an additive structure over certain subsets of items, the auctioneer can find the efficient outcome by using an LP formulation whose dual suggests a pricing rule that also decomposes over these subsets. Moreover, using complementary slackness for LP5G and D5G, it can be established that a pricing equilibrium that relies on a pricing rule with this structure always exists. Thus, when valuations are additively decomposable over a few sets with small cardinality (as in the case of graphical valuations), a simple pricing rule can be used for iterative auction design.

5.6 Summary

In this chapter, we focused on various pricing rules, and investigated the existence of pricing equilibria with these pricing rules. Additionally, we provided LP formulations of the efficient allocation problem that have integral optimal solutions (and can find the efficient allocation), if and only if such equilibria exist. More precisely, we first focused on the anonymous graphical pricing rule. We established that for series-parallel graphs (or graphs that do not contain a 4-clique as a minor) the corresponding LP formulation is equivalent to the anonymous item pricing LP (considered in Chapter 4), in the sense that the former identifies the efficient allocation if and only if the latter does. On the other hand, we showed that if the underlying graph has a 5-clique as a minor, for some choice of weights, this LP formulation can identify the efficient allocation, while anonymous item pricing LP cannot. We then considered LP formulations associated with bidder-specific pricing rules. We showed that in the context of item pricing, allowing for bidder-specific prices is not useful in the sense that the LP formulation associated with bidder-specific pricing rule can identify the efficient outcome, if and only if that associated with anonymous pricing does. On the other hand, the LP formulation associated with bidder-specific graphical prices can find the efficient outcome for all graphical valuations. This result implies that a pricing equilibrium with bidder-specific graphical prices always exist. Hence, efficient iterative auction formats that terminate at such an equilibrium can be developed. We also extended our results to settings where valuation functions admit a general additively decomposable structure, and provided LP formulations that can be used to identify the efficient allocation for such valuation functions. This formulation can be used to develop iterative auction formats that implement the efficient outcome for more general classes of valuation functions.

In Table 5.1, we provide a brief summary of various pricing rules (discussed in this and the previous chapters), and classes of valuations for which the associated LPs can find the

	Anonymous	Bidder-specific
Item pricing	Sign-consistent tree valuations	Equivalent to anonymous item pricing
Graphical pricing	Equivalent to anonymous item pricing for series-parallel graphs	All Graphical Valuations

Table 5.1: A summary of classes of valuation functions and different pricing rules, for which the associated LPs can find the efficient outcome.

efficient outcome. In this table, by “equivalent” we mean that the LP formulation associated with one pricing rule can find the efficient outcome (as an integral optimal solution) if and only if that associated with another can also do so. In Chapter 6, we discuss iterative solutions of these LP formulations, and explain how they can be used to design iterative auction formats that rely on simple pricing rules and implement the efficient outcome for graphical valuations.

5.7 Appendix: Proof of Theorem 5.2.1

Before we prove Theorem 5.2.1, we cover some graph-theoretic preliminaries that will be used in the proof, and establish an auxiliary result.

We start by defining a partition of a graph. Formally, a collection $\pi = \{A^m\}_{m \in \mathcal{A}}$ is a partition if (i) $|A^m| \geq 1$, (ii) $A^m \cap A^k = \emptyset$ for all $m, k \in \mathcal{A}$, (iii) $\cup_m A^m$ gives the set of all nodes. We denote the set of all partitions associated with a graph G as Π . Note that when defining Π , we do not restrict the cardinality of \mathcal{A} . A partition where $|\mathcal{A}| = k$, is referred to as a k -partition of G .

Some of the edges of the graph can be cut by a given partition, i.e., the end points of an edge may belong to different subsets of the partition. We denote the set of edges that are cut by a partition π by $E(\pi)$, i.e., $E(\pi) = \{(i, j) | i, j \notin A^m \text{ for any } m \in \mathcal{M}\}$. We can associate each partition π with an incidence vector $z(\pi)$ that captures the edges that are cut by the given partition, i.e.,

$$z_e(\pi) = \begin{cases} 1 & \text{for } e \in E(\pi) \\ 0 & \text{otherwise.} \end{cases} \quad (5.4)$$

The graph partitioning polytope $P(G)$ is the polytope associated with these incidence vectors. In particular, for a given graph G , $P(G)$ is the convex hull of the incidence vectors, i.e., $P(G) = \text{conv}\{z(\pi) | \pi \in \Pi\}$.

For series-parallel graphs, this polytope admits an alternative characterization. Consider a cycle $C = (V_c, E_c)$ in a given graph. With any edge $e^* \in E_c$, we can associate the following *cycle inequality*,

$$\sum_{e \in E_c - e^*} z_e - z_{e^*} \geq 0. \quad (5.5)$$

The set of all such inequalities defines a polytope,

$$LP(G) = \{z \in \mathbb{R}^{|E|} \mid z \text{ satisfies (5.5) for all cycles } C; 0 \leq z_e \leq 1 \text{ for } e \in E\}. \quad (5.6)$$

An interesting result due to (Chopra, 1994) suggests that for series-parallel graphs

$$P(G) = LP(G), \quad (5.7)$$

i.e., the cycle inequalities fully characterize the partitioning polytope for this class of graphs.

We next establish that for series-parallel graphs, the incidence vectors of partitions can be characterized by restricting attention to 3-partitions.

Lemma 5.7.1. *Let G be a series-parallel graph, and z^* be an incidence vector associated with a k -partition of G , where $k > 3$. Then, there exists a 3-partition of G with the same incidence vector.*

Proof. Consider a given partition $\{S^m\}$ of G with more than 3 sets, and the associated incidence vector z^* . Obtain another graph \hat{G} by contracting every connected subset of nodes of G that belong to the same S^m (for some m). For each node \hat{i} of \hat{G} , denote the set of nodes of G that are contracted to \hat{i} by $\beta(\hat{i})$. Note that \hat{G} is a minor of G . Since G is a series-parallel graph, and series-parallel graphs are equivalent to graphs that do not include 4-clique as a minor, it follows that \hat{G} does not have a 4-clique as a minor, and is a series-parallel graph. Series-parallel graphs have a chromatic number of at most three (Jensen and Toft, 1995). Thus it follows that nodes of \hat{G} admit a 3-coloring. Consider a 3-coloring of \hat{G} , and the associated 3-partition $\{\hat{S}^1, \hat{S}^2, \hat{S}^3\}$ of \hat{G} where nodes of same color are assigned to the same set. A corresponding partition $\{\bar{S}^1, \bar{S}^2, \bar{S}^3\}$ of G can be obtained by defining $\bar{S}^k = \cup_{j \in \hat{S}^k} \beta(j)$ for $k \in \{1, 2, 3\}$. We claim that the incidence vector associated with this partition is also z^* . To see this note that if $z_{ij}^* = 1$, then it should be the case that i and j belong to different sets in the original partition $\{S^m\}$. Consequently, in \hat{G} they correspond to adjacent nodes. Thus in the coloring, they have different colors, and hence, they belong to different sets of partition $\{\hat{S}^m\}$. Thus, in the final partition $\{\bar{S}^k\}$ the nodes i and j belong to different sets, and hence the incidence vector associated with this partition satisfies $z_{ij} = 1$. Conversely, if $z_{ij}^* = 0$, then nodes i and j belong to the same set in the original partition $\{S^m\}$. Consequently, in \hat{G} they are represented by the same node and in the final partition $\{\bar{S}^k\}$ they belong to the same set. Thus, the incidence vector associated with $\{\bar{S}^k\}$ also satisfies $z_{ij} = 0$. Therefore, we conclude that it is possible to construct a 3-partition of G that has the same incidence vector z^* as the original partition, and the claim follows. \square

Using Lemma 5.7.1 we next establish Theorem 5.2.1.

Proof of Theorem 5.2.1 (i). Assume that LP1 has an optimal solution that is integral. Consider the associated feasible allocation, whose welfare is equal to the optimal objective value

of LP1. Lemma 5.2.1 (ii) implies that in LP3 there exists a feasible solution that corresponds to this allocation. Moreover, the objective value associated with this solution is equal to the total welfare that can be raised by the aforementioned allocation. Note that this solution is optimal in LP3, since Lemma 5.2.1 (i) implies that no feasible solution of LP3 can have larger objective value than the optimal solution of LP1. Thus, it follows that LP3 also has an optimal integral solution.

Conversely assume that LP3 has an optimal solution that is integral. We next show that there exists an optimal solution in LP1 that is integral.

First, observe that if LP3 has an optimal solution that is integral, then due to Assumption 4.2.1 it has another optimal solution that is integral and satisfies $\sum_m \sum_{S|i \in S} x^m(S) = 1$ for all nodes i . This suggests that if LP3 has an optimal integral solution, then the following LP also has an optimal integral solution, with the same objective value:

$$\begin{aligned}
& \max \sum_m \sum_S x^m(S) v^m(S) \\
& \text{s.t.} \sum_S x^m(S) \leq 1 \quad \forall m \\
& \sum_m \sum_{S|i \in S} x^m(S) = 1 \quad \forall i \\
LP3b : & \sum_m \sum_{S|ij \in S} x^m(S) = \sum_m \sum_{\mu \in \chi | ij \in \mu^m} \delta_\mu \quad \forall ij \in E \\
& \sum_{\mu \in \chi} \delta_\mu = 1 \\
& \delta_\mu \geq 0, x^m(S) \geq 0
\end{aligned}$$

Thus, in order to prove the claim it suffices to prove that if LP3b has optimal integral solutions, then so does LP1.

We next show that it is possible to reformulate LP3b, in terms of the partition polytope associated with the underlying value graph. This new formulation does not rely on the $\{\delta_\mu\}$ variables, and allows us to exploit the properties of the partition polytope (in particular (5.7)) for the proof. In particular, we focus on LP3c given below:

$$\begin{aligned}
& \max \sum_m \sum_S x^m(S) v^m(S) \\
& \text{s.t.} \sum_S x^m(S) \leq 1 \quad \forall m \\
LP3c : & \sum_m \sum_{S|i \in S} x^m(S) = 1 \quad \forall i \\
& \{1 - \sum_m \sum_{S|ij \in S} x^m(S)\}_{ij} \in P(G) \\
& x^m(S) \geq 0.
\end{aligned}$$

Our next lemma establishes the equivalence of LP3c and LP3b.

Lemma 5.7.2. *Assume that there are at least three bidders. Then, LP3b and LP3c equivalent, i.e., if $\{x^m(S)\}$ is feasible in LP3c, then for some $\{\delta_\mu\}$, $\{x^m(S), \delta_\mu\}$ is feasible in LP3b, and conversely if $\{x^m(S), \delta_\mu\}$ is feasible in LP3b, then $\{x^m(S)\}$ is feasible in LP3c.*

Proof. Consider any feasible solution of LP3b. We claim that

$$\{h_{ij}\}_{ij \in E} = \left\{ 1 - \sum_m \sum_{\mu \in \chi | ij \in \mu^m} \delta_\mu \right\}_{ij \in E},$$

belongs to $P(G)$. To see this observe that a complete feasible allocation μ induces a partition of items to k different sets, where k is the number of nonempty components of $\{\mu^m\}$. Assume that $\delta_\mu = 1$ for some μ . The definition of h_{ij} suggests that $h_{ij} = 0$ if $i, j \in \mu^m$ for some m , and $h_{ij} = 1$ if $i \in \mu^m, j \in \mu^k$. This suggests that $\{h_{ij}\}$ is an incidence vector associated with μ . For arbitrary $\{\delta_\mu\}$, since $\sum_\mu \delta_\mu = 1$, it follows that $\{h_{ij}\}_{ij}$ belongs to the convex hull of these incidence vectors, and hence is in $P(G)$. This suggests that if $\{x^m(S), \delta_\mu\}$ is feasible in LP3b, then $\{x^m(S)\}$ is feasible in LP3c.

Conversely, consider some $z \in P(G)$. Observe that by definition of $P(G)$, we have $z = \sum_k \alpha_k z^k$, where $\alpha_k \in [0, 1]$, $\sum_k \alpha_k = 1$, and each z^k is an incidence vector that corresponds to a partition of G . Lemma 5.7.1 implies that each z^k can be associated with a 3-partition. Since there are at least three players, this observation implies that for some choice of $\{\delta_\mu\}$, $z_{ij}^k = 1 - \sum_m \sum_{\mu \in \chi | ij \in \mu^m} \delta_\mu$. Since this is true for any z^k , and z is a convex combination of $\{z^k\}$, it follows that for some choice of δ_μ , we have

$$z_{ij} = 1 - \sum_m \sum_{\mu \in \chi | ij \in \mu^m} \delta_\mu, \quad \text{for all } ij. \quad (5.8)$$

Consider a feasible solution $\{x^m(S)\}$ of LP3c. It follows from (5.8) that there exists some $\{\delta_\mu\}$ such that $\{1 - \sum_m \sum_{S | ij \in S} x^m(S)\}_{ij} = 1 - \sum_m \sum_{\mu \in \chi | ij \in \mu^m} \delta_\mu$ for all edges (i, j) . This implies that $\{x^m(S), \delta_\mu\}$ is feasible in LP3b.

These observations suggest that for any feasible solution $\{x^m(S)\}$ of LP3c we have a feasible solution $\{x^m(S), \delta_\mu\}$ of LP3b and vice versa. Thus, the claim follows. \square

This lemma implies that the projection of the feasible sets of these optimization problems onto $\{x^m(S)\}$ variables coincide. Since the objective is only a function of $\{x^m(S)\}$, it also suggests that they have the same objective value. Additionally, the lemma suggests that in order to prove the claim it is sufficient to show that if LP3c has an optimal solution that is integral, then so does LP1. We next establish this, by exploiting the properties of the partition polytope.

It follows from (5.7) that when G is a series-parallel graph, we can replace the constraint $\{1 - \sum_m \sum_{S | ij \in S} x^m(S)\}_{ij} \in P(G)$ with $\{1 - \sum_m \sum_{S | ij \in S} x^m(S)\}_{ij} \in LP(G)$. The

next lemma suggests that the latter constraint immediately follows from the constraints $\sum_S x^m(S) \leq 1$ and $\sum_m \sum_{S|i \in S} x^m(S) = 1$ in LP3c, and hence can be omitted.

Lemma 5.7.3. *Assume that $\sum_S x^m(S) \leq 1$ and $\sum_m \sum_{S|i \in S} x^m(S) = 1$ for all i . Then, $\{1 - \sum_m \sum_{S|ij \in S} x^m(S)\}_{ij} \in LP(G)$.*

Proof. Note that it immediately follows from $\sum_m \sum_{S|i \in S} x^m(S) = 1$ that

$$1 - \sum_m \sum_{S|ij \in S} x^m(S) \geq 1 - \sum_m \sum_{S|i \in S} x^m(S) \geq 0.$$

That is $x_e \in [0, 1]$ for all edges $e \in E$. Thus, in order to establish the claim it suffices to establish that $\{1 - \sum_m \sum_{S|ij \in S} x^m(S)\}_{ij}$ satisfies the cycle inequalities (5.5).

Consider any cycle C and $e^* \in C$. The cycle constraint associated with $\{1 - \sum_m \sum_{S|ij \in S} x^m(S)\}_{ij}$ is given below:

$$\sum_{e \in E_C - e^*} \left(1 - \sum_m \sum_{S|e \in S} x^m(S) \right) \geq 1 - \sum_m \sum_{S|e^* \in S} x^m(S).$$

Rearranging the terms, we need to show

$$|C| - 2 \geq \sum_{e \in E_C - e^*} \sum_m \sum_{S|e \in S} x^m(S) - \sum_m \sum_{S|e^* \in S} x^m(S). \quad (5.9)$$

Let $\sum_m \sum_{S|e^* \in S} x^m(S) = b$ for some real number b , and denote by e^* the edge (i, j) . Note that

$$\sum_m \sum_{S|i \in S} x^m(S) = 1 = \sum_m \sum_{S|e^* \in S} x^m(S) + \sum_m \sum_{S|i \in S, j \notin S} x^m(S). \quad (5.10)$$

This suggests that

$$\sum_m \sum_{S|i \in S, j \notin S} x^m(S) = 1 - b. \quad (5.11)$$

On the other hand, a similar expression to (5.10) can be obtained for edge $(i', j') \in E_C$, implying $\sum_m \sum_{S|(i', j') \in S} x^m(S) = 1 - \sum_m \sum_{S|i' \in S, j' \notin S} x^m(S)$. Thus, we have

$$\sum_{e \in E_C - e^*} \sum_m \sum_{S|e \in S} x^m(S) = |C| - 1 - \sum_{(i', j') \in E_C - e^*} \sum_m \sum_{S|i' \in S, j' \notin S} x^m(S). \quad (5.12)$$

Observe that

$$\sum_{(i', j') \in E_C - e^*} \sum_m \sum_{S|i' \in S, j' \notin S} x^m(S) \geq \sum_m \sum_{S|i \in S, j \notin S} x^m(S) = 1 - b. \quad (5.13)$$

Here the equality follows from (5.11). On the other hand, the inequality follows from the fact that any set that contains i but not j , cuts the cycle at edge (i, j) and at another edge,

i.e., if $i \in S$, $j \notin S$, then there exists an edge $(i', j') \in E_c$ such that $(i, j) \neq (i', j')$, and $i' \in S$, $j' \notin S$. Together with (5.12) this implies that

$$\sum_{e \in E_c - e^*} \sum_m \sum_{S|e \in S} x^m(S) \leq |C| - 1 - (1 - b).$$

Hence, we conclude

$$\sum_{e \in E_c - e^*} \sum_m \sum_{S|e \in S} x^m(S) - \sum_m \sum_{S|e^* \in S} x^m(S) \leq |C| - 1 - (1 - b) - b = |C| - 2. \quad (5.14)$$

Thus, (5.9) holds, and the claim follows. \square

Thus, the constraint $\{1 - \sum_m \sum_{S|ij \in S} x^m(S)\}_{ij} \in P(G)$ or equivalently (for series-parallel graphs) $\{1 - \sum_m \sum_{S|ij \in S} x^m(S)\}_{ij} \in LP(G)$ in LP3c can be omitted. After omitting the aforementioned constraint, LP3c becomes equivalent to LP1.

Summarizing if LP3 has integral optimal solutions, then LP3b has optimal solutions that are integral as well. This optimization formulation, on the other hand, is equivalent to LP3c. Thus, it follows that LP3c also has optimal solutions that are integral. On the other hand, for series-parallel graphs, Lemma 5.7.3 suggests that some constraints of LP3c can be omitted, without affecting the optimality of this solution. However, after omitting these constraints we observe that LP3c is equivalent to LP1, and hence the integral solution we have is also optimal in LP1. Therefore, we conclude that if LP3 has optimal solutions that are integral, then so does LP1, and the claim follows. \square

Proof of Theorem 5.2.1 (ii). We establish the result by showing that Example 5.2.1 can be “embedded” into any graph that includes K_5 as a minor. The proof follows by an explicit construction of weights for a given graph.

Assume that G has K_5 as a minor. This implies that initially some nodes and edges of G can be deleted, then some of the remaining edges can be contracted to obtain K_5 from G (Diestel, 2005).

We assume that there are three bidders, and all bidders have same weights at the nodes/edges of the underlying graph G . We assign these weights as follows:

- Every node that is initially deleted has weight zero. Additionally, edges that are adjacent to such nodes also have weight zero.
- Every edge that is initially deleted has weight zero.
- Assume that the set of nodes V_i are contracted to a single node in K_5 . All edges (that are not initially deleted) between the nodes in V_i are assigned a large weight L .
- Assume that nodes i and j of K_5 , are obtained via contraction of sets V_i and V_j respectively. Assume that there are K_{ij} edges between V_i and V_j that are not initially

deleted. Assume that in Example 5.2.1 the weight on edge (i, j) is w_{ij} . Assign each of the edges between V_i and V_j weights that are equal to w_{ij}/K_{ij} .

- Set the node weight of one node in V_i equal to 100, and the remaining node weights to 0.

Observe that for sufficiently large L , at the efficient allocation, all nodes in V_i will be assigned to a single bidder, for all i . This implies that an efficient allocation can be obtained by assigning $\{V_i\}$ according to the efficient allocation in Example 5.2.1 to the bidders. Moreover, since all bidders have same weights, by choosing the price variables in D3 equal to node/edge weights we obtain a dual feasible solution with objective value equal to the welfare generated by the efficient assignment. Thus, (as in Example 5.2.1), it follows that LP3 has an integral optimal solution for this problem instance. By construction, the corresponding objective value is 525 (the optimal objective value in Example 5.2.1) + $L \times \sum_i \rho_i$, where ρ_i denotes the number of undeleted edges of G between the nodes in V_i .

On the other hand, there exists a feasible solution of LP1 with higher objective value, obtained by choosing $x^m(S)$ in a similar fashion to Example 5.2.1. In particular, if $x^m(S) > 0$ for some S and player m in this example, we construct a feasible solution $\{\bar{x}^m(S)\}$ to LP1 (formulated over G), by setting $\bar{x}^m(\cup_{i \in S} V_i) = x^m(S)$. Observe that by construction the corresponding objective value in LP1 is equal to 527 (the optimal objective value in Example 5.2.1) + $L \times \sum_i \rho_i$. Since a feasible solution of LP1 has higher objective value than the welfare associated with the efficient allocation, we conclude that this problem does not have optimal solutions that are integral.

Thus, it follows that for the above choice of weights, while LP3 has an integral solution, LP1 does not. Since, G is an arbitrary graph that has K_5 as a minor the result follows. \square

Chapter 6

Iterative Algorithms and Auction Design

6.1 Introduction and Organization

In this chapter, we focus on iterative algorithms that can be used for solving two of the LP formulations, LP1 and LP5, introduced in Chapters 4 and 5. Employing these iterative algorithms, we design new iterative auctions where,

- The auctioneer sets prices for the items she sells,
- Bidders respond to the prices chosen by the auctioneer by revealing their demand,
- The auctioneer either terminates the auction with a final allocation, or updates the prices.

The auction formats we obtain terminate when a Walrasian equilibrium or a pricing equilibrium is identified, and implement the efficient allocation for graphical valuations. Importantly, they do so while relying on simple pricing rules (anonymous item pricing or bidder-specific graphical pricing).

Given LP formulations of the efficient allocation problem, our approach for developing iterative auctions has two main components. First we focus on iterative algorithms (such as the primal-dual algorithm) for solving the linear programming formulations of the efficient allocation problem introduced in the previous chapters, and use them to develop new iterative auction formats. Importantly, since bidders' valuations are private knowledge, the convergence of these iterative auctions to the efficient outcome relies on bidders' revelation of their demand truthfully. Secondly, we show that it is possible to guarantee that bidders reveal their demand truthfully at an ex-post perfect equilibrium of these iterative auctions, by charging final payments to bidders that are related to the prices that emerge at the end of this iterative process. This approach for iterative auction design is also employed in the existing literature (see Vohra (2011)). However, majority of the existing iterative

auction formats rely on exponentially many prices for implementing the efficient outcome (Bikhchandani et al., 2002; De Vries et al., 2007; Ausubel and Milgrom, 2002; Mishra and Parkes, 2007; Vohra, 2011). In contrast, in our work, we follow a similar approach, but by exploiting the properties of graphical valuations, we obtain iterative auction formats that rely on simple pricing rules, and guarantee efficiency for all graphical valuations (including those that exhibit complementarity).

For simplicity, we establish our results under an additional assumption which we impose throughout the chapter: Value functions of bidders are integer-valued.¹ We start presenting our approach, by briefly summarizing primal-dual algorithms in Section 6.2, in a general abstract setup. The results and insights we discuss in this section, are used in subsequent sections to establish convergence of iterative auctions we design to the efficient outcome.

In Section 6.3, we introduce the solution concept we use for analyzing the outcome of iterative auctions. In particular, in this section we discuss the ex-post perfect equilibrium concept. Additionally, we provide a sufficient condition, which can be checked to determine if a given strategy profile is an ex-post perfect equilibrium. The results of this section are applicable in general iterative auction environments, and do not rely on the graphical valuation structure.

In Section 6.4, we focus on sign-consistent tree valuations, and LP1, which can be used to find the efficient outcome for this class of valuation functions. We show that iterative solutions of LP1 using primal-dual algorithms, suggest an iterative auction format where the auctioneer offers an anonymous price for each item she is selling, and players report the bundles that they demand at the prices set by the auctioneer. If an efficient allocation is not found, the auctioneer decreases the prices of the “underdemanded” items, and increases the prices of the “overdemanded” ones. This iterative process converges to a Walrasian equilibrium and can be used to implement the efficient outcome when bidders reveal their demand truthfully at each stage of the iterative auction. In order to guarantee truthful bidding, we propose running a series of iterative auctions. In the serial auction setup, the auctioneer uses the first auction to identify an efficient allocation of items to bidders, and the subsequent auctions to find final payments that ensure that bidders have incentive to truthfully reveal their demand. Our results suggest that by running serial auctions that rely on anonymous item pricing, the efficient outcome can be guaranteed to emerge at an ex-post perfect equilibrium for sign-consistent tree valuations.² Additionally, we discuss the problem of iterative auction design with anonymous item pricing, in settings where the

¹This assumption is commonly made in the context of iterative auctions and it usually leads to auction formats with unit price increments, see for instance (De Vries and Vohra, 2003; Ausubel, 2004; Mishra and Parkes, 2007; Bikhchandani et al., 2011; Vohra, 2011).

²In various iterative auction settings (such as Ausubel (2006)) in order to guarantee truthful bidding by bidders, it is necessary to either run multiple iterative auctions (in series or parallel), or use a payment scheme that potentially leads to negative payoffs for some bidders at the end of the auction (thereby deterring entry to the auctions). While this feature is present in the iterative auctions we discuss in Section 6.4, in Section 6.5 we establish that when the auctioneer relies on bidder-specific graphical pricing, it is possible to implement the efficient outcome by running only a single iterative auction.

auctioneer can run only a single iterative auction.

In Section 6.5, we consider general graphical valuations, and LP5 that can be used to find the efficient outcome for all such valuations. By studying the solution of LP5 using iterative algorithms we obtain iterative auction formats where truthful bids of agents converge to a pricing equilibrium, and implement the efficient outcome. The prices that are employed in the resulting iterative auction are bidder-specific graphical prices. In order to guarantee that bidders participate in this auction truthfully, it is possible to run a series of auctions as we suggest in Section 6.4. On the other hand, we show that when the auctioneer employs bidder-specific graphical prices, there exists a price (or dual variable) update rule that the auctioneer can use to solve LP5, such that at the end of the auction the prices converge to a pricing equilibrium (with bidder-specific graphical prices) that reveals (i) the efficient outcome, and (ii) the final payments which guarantee that bidders truthfully reveal their information at each stage of the iterative auction. This price update rule closely relates to primal-dual algorithms, as we explain in detail in Section 6.5. Using this price update rule for the solution of LP5, eliminates the need to run a series of auctions for implementing the efficient outcome. Consequently, in this section we obtain an iterative auction format that relies on bidder-specific graphical pricing, and implements the efficient outcome for general graphical valuations without running a series of auctions. The iterative auction formats that we obtain in Sections 6.4 and 6.5 can be viewed as multi-item generalizations of single-item efficient iterative auctions (e.g., English and Dutch auctions).

In Section 6.5, we are able to obtain iterative auctions that implement the efficient outcome without employing a series of auctions, since the dual of LP5 (D5) has multiple optimal solutions, some of which can be used to compute final payments that guarantee truthful demand reports. On the other hand, there may be optimal solutions to D5, that do not have this property. Hence, a special price update rule that converges to a particular dual optimal solution is necessary in order to implement the efficient outcome. We provide, in Section 6.6, a similar LP formulation, whose corresponding dual optimal solutions always allow for computation of final payments that guarantee truthful bidding. Thus, solution of this LP formulation with *any* primal-dual algorithm can be used to develop iterative auction formats that implement the efficient allocation. Moreover, in Section 6.6, we show that our final formulation can be generalized to obtain iterative auctions that implement the efficient outcome for valuation functions that are not necessarily graphical. In particular, in this section, we outline a framework that allows for implementing the efficient outcome for general additively decomposable valuations.

We close the chapter in Section 6.7, with a brief summary of its contributions. Some of the proofs are presented in Section 6.8 to simplify the exposition.

6.2 Iterative Algorithms for Linear Programs

In this section, we focus on a generic linear program:

$$\begin{aligned} \max \quad & d^T y \\ \text{s.t.} \quad & Ay \leq b \\ & y \geq 0, \end{aligned} \tag{6.1}$$

where $b \in \mathbb{R}^m$, $d \in \mathbb{R}^n$ and $A \in \mathbb{R}^{m \times n}$ are given, and $y \in \mathbb{R}^n$ is the vector of decision variables for this problem. We denote the element in the i th row and j th column of matrix A by A_{ij} , and the i th element of a given vector y by y_i . Observe that the linear programs that we presented for finding the efficient allocation (in particular LP1 and LP5) are special cases of (6.1).

We next explain how this problem can be iteratively solved by using a primal-dual algorithm. The results of this subsection are already present in the literature (for instance see Papadimitriou and Steiglitz (1998); Bertsimas and Tsitsiklis (1997)), and our presentation here closely follows that of Vohra (2011). The results of this section are used in subsequent sections for developing iterative auction formats that implement the efficient outcome.

The dual of the linear program in (6.1) can be given as follows:

$$\begin{aligned} \min \quad & b^T \lambda \\ \text{s.t.} \quad & A^T \lambda \geq d \\ & \lambda \geq 0, \end{aligned} \tag{6.2}$$

where $\lambda \in \mathbb{R}^m$ is the vector of dual variables.

Assume that λ^* is dual feasible, y^* is primal feasible, and these vectors satisfy complementary slackness conditions, i.e.,

$$\begin{aligned} \lambda_i^* > 0 &\rightarrow \sum_j A_{ij} y_j^* = b_j \\ y_j^* > 0 &\rightarrow \sum_i A_{ij} \lambda_i^* = d_j. \end{aligned} \tag{6.3}$$

Then, it follows that y^* is optimal for the primal problem, and λ^* is optimal for the dual problem (Bertsimas and Tsitsiklis, 1997).

The high-level idea behind primal-dual algorithms is to start with a dual feasible solution, and check if there exists a primal feasible solution that satisfies the complementary slackness condition with the given dual solution. If this is the case, optimal solutions to both problems are found, otherwise, another dual feasible solution with improved objective value can be obtained. The primal-dual algorithm iteratively updates the dual solution, when there is no primal feasible solution satisfying the complementary slackness condition

with the current dual solution, and terminates when such a primal feasible solution can be found.

More precisely, let λ^* denote the dual feasible solution at a given iteration. In order to see if there exists a primal feasible solution satisfying the complementary slackness conditions with λ^* , we consider the following set of constraints:

$$\begin{aligned} \sum_j A_{ij}y_j &= b_i && \text{for all } i \text{ such that } \lambda_i^* > 0 \\ y_j &= 0 && \text{for all } j \text{ such that } \sum_i A_{ij}\lambda_i^* > d_j \\ Ay &\leq b \\ y &\geq 0. \end{aligned} \tag{6.4}$$

The first two constraints in the above system are obtained from (6.3) (by taking the complement of the second statement), and the last two are the primal feasibility constraints. If there exists a y^* satisfying this system, then λ^* and y^* are respectively dual and primal optimal.

If this is not the case, we conclude that there does not exist $y \geq 0, z \geq 0$ such that

$$\begin{aligned} \sum_j A_{ij}y_j &= b_i && \text{for all } i \text{ such that } \lambda_i^* > 0 \\ y_j &= 0 && \text{for all } j \text{ such that } \sum_i A_{ij}\lambda_i^* > d_j \\ \sum_j A_{ij}y_j + z_i &= b_i && \text{for all } i \text{ such that } \lambda_i^* = 0 \\ y, z &\geq 0, \end{aligned} \tag{6.5}$$

where z is a vector of slack variables. In this case, the Farkas lemma implies that there exists some $\bar{\lambda}, \mu \in \mathbb{R}^m$ such that

- $A^T\bar{\lambda} + \mu \geq 0$, where $\mu_j = 0$ except for j satisfying $\sum_i A_{ij}\lambda_i^* > d_j$,
- $\bar{\lambda}_i \geq 0$ for all i satisfying $\lambda_i^* = 0$, and
- $\bar{\lambda}^T b < 0$.

Let $\epsilon > 0$ be a small enough constant. Consider updating the original dual feasible solution λ^* to $\lambda^* + \epsilon\bar{\lambda}$. Here, we refer to $\bar{\lambda}$ as the dual update direction, and $\epsilon > 0$ as the dual update step size. Observe that for j satisfying $\sum_i A_{ij}\lambda_i^* > d_j$, we have $\sum_i A_{ij}(\lambda_i^* + \epsilon\bar{\lambda}_i) > d_j$. On the other hand, for j satisfying $\sum_i A_{ij}\lambda_i^* = d_j$, $\mu_j = 0$, and hence the equation $A^T\bar{\lambda} + \mu \geq 0$ implies that $\sum_i A_{ij}\bar{\lambda}_i \geq 0$. This suggests that $\sum_i A_{ij}(\lambda_i^* + \epsilon\bar{\lambda}_i) \geq d_j$ for such j as well. Additionally, since $\bar{\lambda}_i \geq 0$ for all i satisfying $\lambda_i^* = 0$, it also follows that $\lambda^* + \epsilon\bar{\lambda} \geq 0$. Thus, it follows that $\lambda^* + \epsilon\bar{\lambda}$ is dual feasible. Moreover, since $b^T\bar{\lambda} < 0$, we obtain

$$b^T(\lambda^* + \epsilon\bar{\lambda}) = b^T\lambda^* + \epsilon b^T\bar{\lambda} < b^T\lambda^*.$$

That is $\lambda^* + \epsilon \bar{\lambda}$ is a dual feasible solution with strictly better objective value than λ^* .

Summarizing, for a given dual feasible solution λ^* , either a solution of (6.4) exists and λ^* together with this solution yield optimal primal and dual solutions, or we can construct another dual feasible solution $\lambda^* + \epsilon \bar{\lambda}$ with strictly better cost (where $\epsilon > 0$, and $b^T \bar{\lambda} < 0$). Updating the dual solution to such a feasible solution (whenever a solution to (6.4) does not exist), this process converges to optimal primal and dual solutions. Moreover, choosing ϵ appropriately, convergence takes place in finite time (Papadimitriou and Steiglitz, 1998). We refer to such iterative algorithms as primal-dual algorithms. We note that in primal-dual algorithms there are multiple ways of choosing the update direction $\bar{\lambda}$ for the dual feasible solution.³

Applications of primal-dual algorithms to auction design appeared in the existing literature Vohra (2011), De Vries et al. (2007), and Bikhchandani et al. (2002). In these works, authors model the problem of finding the efficient allocation as a linear program, in settings where (i) valuations do not exhibit a special structure, or (ii) they satisfy a gross substitutes condition (and hence do not allow for value complementarities). In these settings, they develop iterative auction formats by considering iterative solutions of these linear programs via iterative algorithms, such as primal-dual methods. The auction formats obtained in the first setting rely on complex bundle pricing rules (where the auctioneer offers a price for each bundle of items), whereas those in the second setting, do not guarantee efficiency when valuations exhibit complementarities. In subsequent sections, we use primal-dual algorithms to develop new iterative auction formats that guarantee efficiency in settings where valuations of bidders can exhibit complementarities. In doing so, we exploit the structure of graphical valuations, and the associated LP formulations (LP1 and LP5) developed in Chapters 4 and 5. This enables us to have iterative auction formats that rely on simpler pricing rules than the bundle pricing rule.

6.3 Ex-post Perfect Equilibrium

In this section, we introduce the solution concept that we use in our study of iterative auctions. In particular, we discuss ex-post perfect equilibrium. We also provide sufficient conditions for characterizing the ex-post perfect equilibria of iterative auctions. We start by providing some necessary notation and definitions.

In iterative auctions, bidders participate in a multi-stage game. Let H_t denote the history of the bidding process until time instant t . We denote the (behavior) strategy of player m whose valuation is v^m by $s^m(v^m)$. This strategy maps every history to an action. More precisely, with a slight abuse of notation, we denote by $s^m(H_t|v^m) \in \Sigma^m(H_t)$ the action player m , whose valuation is v^m , uses at time t , after observing history H_t . Here, $\Sigma^m(H_t)$ denotes the set of actions player m can use after history H_t .

³A systematic approach for obtaining a direction involves formulating the “restricted primal” problem, and considering its dual. See Papadimitriou and Steiglitz (1998), and Vohra (2011).

Consider a strategy profile $\{s^m\}_m$. For a realization of valuations of bidders $\{v^k\}_k$, the payoff bidder m receives at the end of the auction game is denoted by $u^m(s^m(v^m), s^{-m}(v^{-m})|v^m)$, where $s^{-m}(v^{-m})$ denotes the strategies of all players but m . Similarly, denote by

$$u^m(s^m(v^m), s^{-m}(v^{-m})|H_t, v^m),$$

the payoff bidder m , who is of type v^m , receives by using strategy $s^m(v^m)$, after history H_t , given her opponents use strategies $s^{-m}(v^{-m})$.

Using this notation we next define two closely related solution concepts:

Definition 6.3.1 (Ex-post equilibrium). *A strategy profile $s = \{s^m\}$ is an ex-post equilibrium if it satisfies*

$$u^m(s^m(v^m), s^{-m}(v^{-m})|v^m) \geq u^m(z^m, s^{-m}(v^{-m})|v^m). \quad (6.6)$$

for any realization $\{v^m\}$ of valuations of bidders, player m , and strategy z^m .

This definition suggests that a strategy profile is an ex-post equilibrium, if for any realization of types (or valuations $\{v^m\}$), given strategies of her opponents, no agent has incentive to deviate. In other words, the given strategy profile is a Nash equilibrium of the game, where types are public knowledge. The second solution concept is a refinement of ex-post equilibrium that ensures that players have no incentive to deviate from their strategy, after any realization of the history.

Definition 6.3.2 (Ex-post perfect equilibrium). *A strategy profile $s = \{s^m\}$ is an ex-post perfect equilibrium if after any history H_t , it satisfies*

$$u^m(s^m(v^m), s^{-m}(v^{-m})|H_t, v^m) \geq u^m(z^m, s^{-m}(v^{-m})|H_t, v^m), \quad (6.7)$$

for any realization $\{v^m\}$ of valuations of bidders, player m , and any strategy z^m .

This definition suggests that a strategy profile is an ex-post perfect equilibrium, if after any history, for any realization of valuations, given strategies of her opponents, no agent has incentive to deviate from her strategy. In other words, after any realization of the history and payoffs, the given strategy profile remains to be a Nash equilibrium of the induced subgame, where types of agents are public knowledge.

We next provide a sufficient condition for a strategy profile to be an ex-post perfect equilibrium in an iterative auction game. Before we state our result, we need one more definition (Nisan et al., 2007):

Definition 6.3.3 (VCG Mechanism). *Consider a collection of valuation functions $\{v^m\}$. A mechanism is called a VCG (Vickrey - Clarke - Groves) mechanism if it*

- chooses an efficient allocation, i.e., $\{s^m\} \in \arg \max_{\{Z^m\} | Z^k \cap Z^l = \emptyset} \sum_m v^m(Z^m)$

- assigns each agent m a payment

$$\gamma^m(\{S^m\}, \{v^k\}_{k \neq m}) = h^m(v^{-m}) - \sum_{k \neq m} v^k(S^k),$$

where h^m is any real-valued function.

If h^m is such that

$$h^m(v^{-m}) = \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset} \sum_{k \neq m} v^k(Z^k),$$

then, we say that payments of bidders $(\gamma^m(\{S^m\}, \{v^k\}_{k \neq m})) = \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset} \sum_{k \neq m} v^k(Z^k) - \sum_{k \neq m} v^k(S^k)$ are VCG payments with the Clarke pivot rule. Intuitively, for a given agent, the VCG payment with the Clarke pivot rule captures the opportunity cost she creates in the system, i.e., the difference between the maximum welfare that can be achieved by the remaining agents and the total welfare those agents have at the efficient allocation. In this thesis we will only employ VCG payments with the Clarke pivot rule. For simplicity, we refer to these payments as VCG payments. It is known that in sealed bid auctions, charging bidders VCG payments guarantees that the efficient outcome can be implemented in ex-post equilibrium (Nisan et al., 2007; Krishna, 2009).⁴

Consider any realization of valuations $\{v^k\}$, where bidder m receives a set of items S^m in the efficient allocation. Observe that efficiency requires allocating the remaining items to bidders $k \neq m$ according to

$$\arg \max_{\{Z^k\}_{k \neq m} | Z^k \cap Z^l = \emptyset, Z^k \cap S^m = \emptyset} \sum_{k \neq m} v^k(Z^k).$$

This suggests that for any realization $\{v^k\}$ where bidder m receives a set of items S^m in the efficient allocation, her VCG payment can alternatively be expressed as $\gamma^m(\{S^k\}_k, \{v^k\}_{k \neq m}) = \hat{\gamma}^m(S^m, \{v^k\}_{k \neq m})$, where

$$\hat{\gamma}^m(S^m, \{v^k\}_{k \neq m}) \triangleq \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset} \sum_{k \neq m} v^k(Z^k) - \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset, Z^k \cap S^m = \emptyset} \sum_{k \neq m} v^k(Z^k).$$

Thus, under the Clarke pivot rule the VCG payment of an agent can be expressed only as a function of her opponents' valuations, and the set of items she acquires in the auction. We make this dependence explicit using the function $\hat{\gamma}^m$.

The main result of this section makes use of the VCG payments to obtain a sufficient condition for a strategy profile to be an ex-post perfect equilibrium in an iterative auction.

Theorem 6.3.1. *Consider a strategy profile $\{s^k\}$, and let $\{\hat{v}^k\}$ denote the realization of payoffs of bidders. Assume that after any history H_t , if bidders $k \neq m$ follow strategies*

⁴In fact, under this payment rule, in sealed bid auctions it is also a dominant strategy equilibrium for bidders to bid truthfully, as explained in the aforementioned references.

$\{s^k(\hat{v}^k)\}_{k \neq m}$, and

- If bidder m uses strategy $s^m(\hat{v}^m)$, then the auction terminates with the efficient allocation.
- If bidder m uses a strategy z^m (possibly identical to $s^m(\hat{v}^m)$), such that the auction terminates with bidder m receiving a set of items S^m , then the corresponding payment of bidder m is equal to $\hat{\gamma}^m(S^m, \{\hat{v}^k\}_{k \neq m})$.
- If the auction does not terminate, then bidders do not make payments or receive items.

The strategy profile $\{s^m\}$ is an ex-post perfect equilibrium.

Proof. Let $\{\hat{S}^k\}$ denote the efficient allocation associated with valuations $\{\hat{v}^k\}$. From the first and second conditions of the theorem it follows that for agent m , the payoff the strategy profile $\{s^k(\hat{v}^k)\}$ leads to (after history H_t), is given by

$$\begin{aligned}
u^m(s^m(\hat{v}^m), s^m(\hat{v}^{-m}) | H_t, \hat{v}^m) &= \hat{v}^m(\hat{S}^m) - \hat{\gamma}^m(\hat{S}^m, \{\hat{v}^k\}_{k \neq m}) \\
&= \hat{v}^m(\hat{S}^m) + \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset, Z^k \cap \hat{S}^m = \emptyset} \sum_{k \neq m} \hat{v}^k(Z^k) - \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset} \sum_{k \neq m} \hat{v}^k(Z^k) \\
&= \sum_k \hat{v}^k(\hat{S}^k) - \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset} \sum_{k \neq m} \hat{v}^k(Z^k).
\end{aligned} \tag{6.8}$$

Since $\{\hat{S}^k\}$ is the efficient allocation we have $\sum_k \hat{v}^k(\hat{S}^k) - \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset} \sum_{k \neq m} \hat{v}^k(Z^k) \geq 0$, and hence (6.8) implies that $u^m(s^m(\hat{v}^m), s^m(\hat{v}^{-m}) | H_t, \hat{v}^m) \geq 0$. Thus, after history H_t , bidder m has no incentive to use a strategy that prevents the auction from terminating.

Assume that after history H_t bidder m can use a strategy z^m so that the auction terminates with allocation $\{S^m\}$. The second condition of the theorem together with (6.8) implies that

$$\begin{aligned}
u^m(s^m(\hat{v}^m), s^m(\hat{v}^{-m}) | H_t, \hat{v}^m) &= \sum_k \hat{v}^k(\hat{S}^k) - \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset} \sum_{k \neq m} \hat{v}^k(Z^k) \\
&\geq \hat{v}^m(S^m) + \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset, Z^k \cap S^m = \emptyset} \sum_{k \neq m} \hat{v}^k(Z^k) - \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset} \sum_{k \neq m} \hat{v}^k(Z^k) \\
&= \hat{v}^m(S^m) - \hat{\gamma}^m(S^m, \{\hat{v}^k\}_{k \neq m}) = u^m(z^m, s^m(\hat{v}^{-m}) | H_t, \hat{v}^m),
\end{aligned}$$

where the inequality follows from the fact that $\{\hat{S}^k\}$ is the efficient allocation, and the last equality follows from the second condition of the theorem. Thus, we conclude that bidder m cannot improve her payoff by deviating from $\hat{s}^m(\hat{v}^m)$ to another strategy that leads the auction to terminate. Since bidder m , history H_t , strategy z^m , and payoff realization $\{\hat{v}^k\}$ are arbitrary, the claim follows from Definition 6.3.2. \square

This theorem suggests that if a strategy profile leads to the efficient outcome, and is such that (i) when her opponents follow their strategies, the payment of an agent is only a

function of her opponents' true valuations, and the final bundle of items she receives, and (ii) corresponds to VCG payments specified by this bundle and the opponents' valuations, then it is an ex-post perfect equilibrium. In Sections 6.5, we propose iterative auction formats where the conditions of the above theorem are satisfied, and the efficient outcome emerges at an ex-post perfect equilibrium.

6.4 Iterative Auctions for Tree Valuations

Corollary 4.3.1 suggests that an optimal solution of the optimization problem LP1 of Chapter 4 provides the efficient allocation when the underlying valuation functions satisfy Assumptions 4.3.1 and 4.3.2, i.e., when the valuations are sign-consistent tree valuations. In this section, we first explain how the primal-dual algorithm presented in Section 6.2, can be used to solve this optimization problem (Section 6.4.1). We then use the resulting algorithm to design an iterative auction format that implements the efficient outcome for sign-consistent tree valuations (Section 6.4.2). In order to simplify exposition, throughout this section, we assume that valuation functions are integer-valued.

6.4.1 An Iterative Algorithm for LP1

Recall that LP1 and its dual, which were introduced in Chapter 4, are as given below:

$$\begin{array}{l|l}
 \begin{array}{l}
 \max \sum_m \sum_S x^m(S) v^m(S) \\
 \text{s.t. } \sum_S x^m(S) \leq 1 \quad \forall m \\
 \sum_m \sum_{S|i \in S} x^m(S) \leq 1 \quad \forall i \\
 x^m(S) \geq 0.
 \end{array} & \begin{array}{l}
 \min \sum_i p_i + \sum_m \pi^m \\
 \text{s.t. } \pi^m \geq v^m(S) - \sum_{i \in S} p_i \quad \forall S, m \\
 p_i, \pi^m \geq 0.
 \end{array} \\
 (LP1) & (D1)
 \end{array}$$

In this section, we provide a primal-dual algorithm for solutions of these LPs. We adopt the shorthand notations π and p to denote the vectors of $\{\pi^m\}_m$ and $\{p_i\}_i$ of variables.

The primal-dual algorithm, presented in Section 6.2, involves a dual feasible solution at each step of the algorithm. Observe that at an optimal solution of D1, the dual variable π^m is such that $\pi^m = \max_{S \subset \mathcal{N}} v^m(S) - \sum_{i \in S} p_i$ for all $m \in \mathcal{M}$. Motivated by this observation, at each step of our primal-dual algorithm, we will choose the π vector so that $\pi^m = \max_{S \subset \mathcal{N}} v^m(S) - \sum_{i \in S} p_i$ for all $m \in \mathcal{M}$. Observe that this construction satisfies the first constraint of D1 trivially.

As explained in Chapter 4, the variable p_i can be thought of as the price of item i , and π^m can be thought of as the maximum surplus of player m , at the given prices. Also recall that we say that a bundle S^* is *demanded* by player m if $\pi^m = v^m(S^*) - \sum_{i \in S^*} p_i$, and the set of items a bidder demands is denoted by D^m , i.e., $D^m = \{S | \pi^m = v^m(S) - \sum_{i \in S} p_i\}$.

Given a dual feasible solution (p, π) of D1, the primal-dual algorithm first checks if a primal feasible solution that satisfies the complementary slackness conditions with (p, π) exists. In particular, following (6.4), we check whether a solution $\{x^m(S)\}$ to the following system of equations exist:

$$\begin{aligned}
\sum_S x^m(S) &= 1 && \text{for all } m \text{ such that } \pi^m > 0, \\
\sum_m \sum_{S|i \in S} x^m(S) &= 1 && \text{for all } i \text{ such that } p_i > 0, \\
x^m(S) &= 0 && \text{for all } m, S \text{ such that } S \notin D^m, \\
\sum_S x^m(S) &\leq 1 && \text{for all } m \text{ such that } \pi^m = 0, \\
\sum_m \sum_{S|i \in S} x^m(S) &\leq 1 && \text{for all } i \text{ such that } p_i = 0, \\
x^m(S) &\geq 0 && \text{for all } m, S \text{ such that } S \in D^m.
\end{aligned} \tag{6.9}$$

As explained earlier, if this system of equations has a solution, then it is an optimal solution of the primal problem LP1. Observe that in this formulation, we omit the primal feasibility constraints $\sum_S x^m(S) \leq 1$ for m such that $\pi^m > 0$, and $\sum_m \sum_{S|i \in S} x^m(S) \leq 1$ for i such that $p_i = 0$, since these constraints are implied by the equality constraints in (6.9) for such i and m (see the first two lines of (6.9)).

By exploiting the structure of the optimization problems LP1 and D1, it is possible to obtain additional properties for the solution of the system of equations (6.9). The following lemma (which is proved in Section 6.8) states some of these properties.

Lemma 6.4.1. *Let Assumptions 4.3.1 and 4.3.2 hold. Assume that for a feasible solution (p, π) of D1, $\{x^m(S)\}$ is a solution to (6.9). Then, there exists another solution $\{\hat{x}^m(S)\}$ to (6.9), such that (i) $\sum_S \hat{x}^m(S) = 1$ for all m , (ii) $\sum_m \sum_{S|i \in S} \hat{x}^m(S) = 1$ for all i , (iii) $\hat{x}^m(S) \in \{0, 1\}$ for all m, S .*

This lemma suggests that at a given dual feasible solution, if there exists a feasible solution of (6.9), then there also exists a feasible integral solution such that $\sum_S \hat{x}^m(S) = 1$ for all m . Consider such a solution of (6.9) (assuming it exists). Observe that at this solution, every bidder m has $x^m(S) = 1$, for a bundle she demands (since by the third constraint $x^m(S) = 0$ for bundles that are not demanded), and the resulting allocation is complete and feasible. This implies that finding an integral solution of (6.9) (and hence an efficient allocation) is equivalent to finding a complete feasible allocation, which assigns each bidder a set of items that she demands at the given prices.

Assume that a solution of (6.9) does not exist. Then, the primal-dual algorithm suggests updating the dual variables to a dual feasible solution $(p, \pi) + \epsilon(\bar{p}, \bar{\pi})$ using some step size $\epsilon > 0$, and dual update direction $(\bar{p}, \bar{\pi})$ such that $\sum_i \bar{p}_i + \sum_m \bar{\pi}^m < 0$. Existence of such a

dual feasible solution is guaranteed by Farkas lemma, as explained in Section 6.2. We next show how an improvement direction can systematically be obtained.

Consider the following optimization problem:

$$\begin{aligned}
\min \quad & \sum_{i|p_i>0} \gamma_i + \sum_i h_i \\
s.t. \quad & \sum_S x^m(S) = 1 \quad \text{for all } m \\
& \sum_m \sum_{S|i \in S} x^m(S) + \gamma_i - h_i = 1 \quad \text{for all } i \text{ such that } p_i > 0 \\
& \sum_m \sum_{S|i \in S} x^m(S) - h_i \leq 1 \quad \text{for all } i \text{ such that } p_i = 0 \\
& x^m(S) = 0 \quad \text{for all } m, S \text{ such that } S \notin D^m \\
& x^m(S), \gamma_i, h_i \geq 0 \quad \text{for all } m, S, i.
\end{aligned} \tag{6.10}$$

Lemma 6.4.1 suggests that if a solution to (6.9) exists, then there exists a solution for which $\sum_S x^m(S) = 1$ for all m . This implies that if the system in (6.9) has a solution, then the optimal value of (6.10) is zero. Conversely, if (6.9) does not have a solution, then the objective value of the above problem is strictly positive (otherwise an optimal solution of (6.10) gives a solution of (6.9)). Intuitively, the parameters γ_i and h_i measure how much the constraints for item i present in (6.9) are violated, and the objective of the above problem is to minimize the total violation of the constraints. This problem is sometimes referred to as the restricted primal problem (Papadimitriou and Steiglitz, 1998).

Note that in (6.10) some $x^m(S)$ variables are set equal to zero (see the fourth constraint), and this problem can be reformulated omitting these variables. The dual of the resulting problem (also known as the restricted dual problem) can be given as follows:

$$\begin{aligned}
\max \quad & - \sum_m \bar{\pi}^m - \sum_i \bar{p}_i \\
s.t. \quad & \bar{\pi}^m + \sum_{i \in S} \bar{p}_i \geq 0 \quad \text{for } m, S \text{ such that } S \in D^m \\
& \bar{p}_i \leq 1 \quad \text{for all } i \\
& -\bar{p}_i \leq 1 \quad \text{for all } i \text{ such that } p_i > 0, \\
& \bar{p}_i \geq 0 \quad \text{for all } i \text{ such that } p_i = 0.
\end{aligned} \tag{6.11}$$

Here $\bar{\pi}^m$ is the dual variable corresponding to the first constraint of (6.10), and \bar{p}_i is the dual variable corresponding to the second and third constraints. In (6.10), γ_i is the variable corresponding to the dual constraint $-\bar{p}_i \leq 1$ and h_i is the variable corresponding to the constraint $\bar{p}_i \leq 1$.

If at an optimal solution of (6.10), $h_i > 0$, then we will say that item i is *overdemanded*. Similarly, if at an optimal solution $\gamma_i > 0$, we will say that i is *underdemanded*. Intuitively,

overdemanded items are demanded by more than one bidder, whereas the converse is true for the underdemanded items. Note that (6.10) may have multiple optimal solutions, but it is never the case that for some item i , $h_i > 0$ in some optimal solutions, and $\gamma_i > 0$ in the others. This is because, when $h_i > 0$ at an optimal solution, the complementary slackness conditions suggest that in the dual optimal solutions the constraint $\bar{p}_i \leq 1$ is active. Conversely if $\gamma_i > 0$ at an optimal solution, then the constraint $\bar{p}_i \geq -1$ is active. The claim follows, since, at a given dual optimal solution, at most one of these inequalities can be active. This suggests that an item can be either overdemanded or underdemanded but not both.

Assume that (6.9) has no feasible solution. Then the optimal value of (6.10) is positive. Additionally, by strong duality, (6.11) has an optimal solution $(\bar{p}, \bar{\pi})$ that leads to a positive objective value that is equal to the optimal value of (6.10).

We next establish that $(\bar{p}, \bar{\pi})$ obtained by solving (6.11), is a *valid dual update direction* for our primal-dual algorithm, i.e., it is such that $\sum_i \bar{p}_i + \sum_m \bar{\pi}^m < 0$ and for small enough $\epsilon > 0$, $(p, \pi) + \epsilon(\bar{p}, \bar{\pi})$ is dual feasible. Moreover, this dual update direction leads to an intuitive price update structure: the auctioneer increases the prices of overdemanded items, and decreases the prices of the underdemanded ones.

Lemma 6.4.2. *Let (p, π) be a dual feasible solution of D1 such that (6.9) has no solution corresponding to it. Consider an optimal solution $(\{x^m(S)\}, \{\gamma_i\}, \{h_i\})$ of (6.10) and an optimal solution $(\bar{p}, \bar{\pi})$ of (6.11). The solution $(\bar{p}, \bar{\pi})$ is a valid dual update direction. Moreover, $(\bar{p}, \bar{\pi})$ satisfies the following conditions:*

- (i) *If i is overdemanded, then $\bar{p}_i > 0$.*
- (ii) *If i is underdemanded then $\bar{p}_i < 0$.*
- (iii) *Assume that $\pi^m = \max_S v^m(S) - \sum_{i \in S} p_i$ for some player m . There exists some set S^* such that $\pi^m = v^m(S^*) - \sum_{i \in S^*} p_i$, and $\pi^m + \epsilon \bar{\pi}^m = v^m(S^*) - \sum_{i \in S^*} (p_i + \epsilon \bar{p}_i)$, for any $\epsilon > 0$ such that $(p, \pi) + \epsilon(\bar{p}, \bar{\pi})$ is feasible in D1.*

This lemma (proof of which can be found in Section 6.8) implies that for solving LP1 using a primal-dual algorithm, a valid dual update direction can be found by solving (6.11). Additionally, it suggests that for each player m there exists some set S that is demanded both at the original prices, and the updated prices. Another implication of the third part of the lemma is that if at the initial dual feasible solution, we have $\pi^m = \max_S v^m(S) - \sum_{i \in S} p_i$, then provided that the dual update direction is computed using (6.11) (and a sufficiently small step size is chosen), this equality holds at all steps of the iterative algorithm.

The only unspecified step in our primal-dual algorithm is the choice of ϵ , the step size of the dual update. In general, in primal-dual algorithms, ϵ is chosen as the largest step size that does not violate the dual feasibility constraint (Papadimitriou and Steiglitz, 1998). In our case this translates to choosing ϵ equal to the largest θ satisfying $p_i + \theta \bar{p}_i \geq 0$, and

$\pi^m + \theta \bar{\pi}^m \geq v^m(S) - \sum_{i \in S} (p_i + \theta \bar{p}_i)$ for all i, m, S . Denote such a θ by θ^* . Observe that computing θ^* requires knowledge of valuation functions. Our next result establishes that when valuation functions are integer-valued, it is possible to choose a step size that is only as a function of p and \bar{p} and that preserves dual feasibility.

Lemma 6.4.3. *Assume valuation functions are integer-valued. There exists some function $\phi : \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}_{++}$, such that for any dual feasible solution (p, π) of D1 that is not optimal, and the corresponding update direction $(\bar{p}, \bar{\pi})$,*

- (i) *A dual feasible solution with improved objective is given by $(\hat{p}, \hat{\pi}) = (p, \pi) + \phi(p, \bar{p})(\bar{p}, \bar{\pi})$.*
- (ii) *Either $\phi(p, \bar{p}) = \theta^*$, or $(\bar{p}, \bar{\pi})$ is an optimal solution of (6.11) at the updated dual solution $(\hat{p}, \hat{\pi})$ as well.*
- (iii) *If $\phi(p, \bar{p}) \neq \theta^*$, after finitely many dual variable updates (using the update direction $(\bar{p}, \bar{\pi})$, and choosing the step size according to ϕ), the dual solution $(p, \pi) + \theta^*(\bar{p}, \bar{\pi})$ is reached.*

A proof of this lemma is provided in Section 6.8. This lemma implies that the primal-dual algorithm can be implemented by choosing a step size that does not depend on the valuation functions.

Using the results obtained in this section, we next state a primal-dual algorithm (see Algorithm 1) for solving LP1. We conclude this section by proving that Algorithm 1 can be used to obtain an optimal solution of LP1 that is integral.

Algorithm 1 A primal-dual algorithm for solving LP1/D1.

- S1: Start at a dual feasible solution, $p_i = 0$ for all i , and $\pi^m = \max_S v^m(S) - \sum_{i \in S} p_i$.
 - S2: Given a dual feasible solution, construct sets $D^m = \{S | \pi^m = v^m(S) - \sum_{i \in S} p_i\}$.
 - S3: Solve (6.10). If this problem has objective value 0, the solution (p, π) is optimal in D1. Moreover, there also exists an optimal solution of (6.10) that is integral. Go to step S5.
Otherwise, go to step S4.
 - S4: Find a dual variable update direction $(\bar{p}, \bar{\pi})$ by solving (6.11). Update dual solution to $(p, \pi) + \phi(p, \bar{p})(\bar{p}, \bar{\pi})$, go to step S2.
 - S5: Terminate returning the allocation of items, suggested by the optimal integral solution of (6.10), and the dual solution (p, π) .
-

Lemma 6.4.4. *Algorithm 1 terminates in finitely many steps with an optimal solution of D1 and an optimal integral solution of LP1.*

Proof. Since at the initial dual feasible solution we have $\pi^m = \max_S v^m(S) - \sum_{i \in S} p_i$, it follows from Lemma 6.4.2 that the dual feasible solution satisfies this equality at each iteration, and thus $D^m \neq \emptyset$.

It can be seen that if the objective value of (6.10) is equal to zero, then the solution satisfies the complementary slackness conditions in (6.9). Moreover, Lemma 6.4.1 suggests that there exists an integral solution of (6.9) as well, which also leads to an objective value of zero at (6.10). Thus, we conclude that if the objective value of (6.10) is equal to zero, then there exists an optimal solution, that is integral, and that satisfies the complementary slackness conditions in (6.9). Since, it satisfies the complementary slackness conditions, it follows that it is optimal in LP1 and the corresponding dual solution is optimal in D1. Thus, Step S5 of the algorithm implies that if the algorithm terminates, the corresponding primal and dual solutions are optimal.

If the objective value of (6.10) is not equal to zero, then as Lemmas 6.4.2 and 6.4.3 suggest the solution in Step S4 is dual feasible, and has a lower objective value. Moreover, Lemma 6.4.3 implies that starting from a dual feasible solution (p, π) , in finitely many steps a dual feasible solution $(p, \pi) + \theta^*(\bar{p}, \bar{\pi})$ is reached. On the other hand, it is known that using step-size θ^* primal-dual algorithms converge in finitely many steps (Papadimitriou and Steiglitz, 1998). Thus, it follows that the algorithm terminates in finitely many steps, and the claim follows. □

6.4.2 An Efficient Iterative Auction

In the previous section, we provided a primal-dual algorithm that can be used for solving LP1/D1. In this section, we employ this algorithm to design iterative auctions that implement the efficient outcome for sign-consistent tree valuations. In particular, we first show that the auctioneer can simply run Algorithm 1 by asking bidders the set of items that they demand at the given prices (i.e., the sets D^m in Algorithm 1), and adjusting prices accordingly. This process can be interpreted as an iterative auction, which we refer to as the one-stage auction game. We show that if bidders truthfully reveal their demand this iterative auction converges to a Walrasian equilibrium and implements the efficient outcome. On the other hand, bidders may have incentive to misreport their demand, if the auctioneer charges final payments to bidders that are equal to the prices that emerge at the end of the auction. In order to resolve this issue, we propose another related auction, which we refer to as the serial auction. This auction relies on running a series of one-stage auctions to find the efficient outcome, as well as the payments that guarantee truthful revelation of bidders' demand at each step of the auctions. We show that this iterative auction implements the efficient outcome at an ex-post perfect equilibrium, for sign-consistent tree valuations.

In Section 6.4.1, Algorithm 1 implicitly makes use of the value functions $\{v^m\}$. In particular, in the first stage of the algorithm, we define $\pi^m = \max_S v^m(S) - \sum_{i \in S} p_i$.

Additionally, in Step S2 of the algorithm we construct sets D^m , which can be used to formulate optimization problems (6.10) and (6.11). Note that Steps S1 and S2 are the only steps of the algorithm, where information about payoffs is necessary.

We claim that in order to run this algorithm, it is not necessary to explicitly use variables $\{\pi^m\}$. In particular, assume that initially π^m is defined as in Step S1. Lemma 6.4.2 suggests that after each dual variable update these variables satisfy $\pi^m = \max_S v^m(S) - \sum_{i \in S} p_i$. Moreover, the algorithm does not make use of these variables except for identifying sets for which the aforementioned equality holds. These observations imply that the algorithm can be run by implicitly defining $\pi^m = \max_S v^m(S) - \sum_{i \in S} p_i$ for all bidders m and bundles S at every stage of the algorithm, and constructing sets $D^m = \arg \max_S v^m(S) - \sum_{i \in S} p_i$. Note that this choice of π^m ensures that $D^m = \{S | \pi^m = \max_S v^m(S) - \sum_{i \in S} p_i\}$, i.e., the set of items each bidder m demands at the current prices.

Thus, it follows that Algorithm 1 can be reformulated by eliminating the π variables, and using the sets $D^m = \arg \max_S v^m(S) - \sum_{i \in S} p_i$ for formulating optimization problems (6.10) and (6.11). This suggests a natural iterative auction format, which we refer to as the *one-stage auction*:

One-stage auction.

- S1: Start with prices $p_i = 0$ for all i .
 - S2: Ask each bidder m the set of items she demands $D^m = \arg \max_S v^m(S) - \sum_{i \in S} p_i$.
 - S3: Solve (6.10). If this problem has objective value 0, then there exists an optimal solution of (6.10) that is integral. Go to step S5. Otherwise, go to step S4.
 - S4: Update the prices using \bar{p} obtained by solving (6.11) to $p + \phi(p, \bar{p})\bar{p}$. Go to step S2.
 - S5: Terminate returning the allocation of items, suggested by the optimal integer solution of (6.10). Assign each bidder m who receives a set of items S^m , a payment that is equal to $\sum_{i \in S^m} p_i$.
-

Observe that in the above iterative process, the auctioneer, who does not know the value functions of bidders, sets prices, and bidders respond to these prices with their demand sets. In response to these demand sets, the auctioneer adjusts the prices, until a final allocation is obtained. Note that the price updates have an intuitive interpretation. The auctioneer increases the prices of overdemanded items, and decreases the prices of the underdemanded items as suggested by Step S4, and Lemma 6.4.2. The auction terminates, if the condition in Step S3 holds, i.e., none of the items is under/over demanded (and hence (6.10) has objective value zero).

We next use Lemma 6.4.4 to show that if bidders truthfully report their demand sets, i.e., reveal sets $D^m = \arg \max_S v^m(S) - \sum_{i \in S} p_i$ at each step of the auction, then the one-stage auction terminates with an efficient allocation. In fact, we establish a stronger result: If a subset \mathcal{M}_0 of bidders reveal their demand truthfully, after some history H_t , and

the auction terminates with an assignment of items $\{S^m\}$, then $\{S^m\}_{m \in \mathcal{M}_0}$ is an efficient assignment of items $\cup_{m \in \mathcal{M}_0} S^m$ to bidders in \mathcal{M}_0 .

Lemma 6.4.5. (i) *After any history H_t , assume that bidders in $\mathcal{M}_0 \subset \mathcal{M}$ reveal their demand truthfully at the one-stage auction game, and the auction terminates with a final allocation $\{S^m\}$. Then $\{S^m\}_{m \in \mathcal{M}_0}$ is an efficient assignment of items $\cup_{m \in \mathcal{M}_0} S^m$ to bidders in \mathcal{M}_0 .*

(ii) *After any history H_t , if all bidders reveal their demand truthfully, the one-stage auction game terminates with an efficient allocation of items.*

A proof of this Lemma can be found in Section 6.8.

Observe that Step S2 of the one-step auction game implies that if bidders truthfully reveal their demand, then the allocation and prices obtained at the end of the iterative auction constitute a Walrasian equilibrium. Thus, we conclude that this iterative auction terminates when a natural market clearance condition holds.

Despite the fact that the one-stage auction game converges to the efficient allocation, when agents reveal their demand truthfully, the prices that emerge at the end of this iterative auction may not guarantee truthful demand revelation. We next introduce an extension of this iterative auction, which we refer to as the serial auction game. In this game, bidders compete not in one but multiple one-stage auctions, conducted in series. At the end of this sequence of auctions, we guarantee that the auctioneer finds the efficient outcome, as well as the payments that guarantee truthful revelation of demand sets at an equilibrium. Before we explain our approach in detail, we introduce some additional notation.

Consider the one-stage auction game introduced above. Denote by $p(t, S)$ the total price of items in bundle S , at step $t \in \mathbb{Z}_{++}$ of this auction, i.e., $p(t, S) = \sum_{i \in S} p_i(t)$, where $p(t)$ is the price vector at step t . Let $D^m(t)$ be the set of items bidder m reports as her demand at time t , in response to prices $p(t)$. For each bidder m , and $t \geq 1$, we define

$$q^m(t) = \begin{cases} \max_{S \in D^m(t) \cap D^m(t+1)} p(t+1, S) - p(t, S) & \text{if } D^m(t) \cap D^m(t+1) \neq \emptyset \\ L & \text{otherwise,} \end{cases} \quad (6.12)$$

where L is a large constant. Observe that if bidders reveal their demand sets truthfully, and $D^m(t) \cap D^m(t+1) \neq \emptyset$, we have $v^m(S) - p(t, S) = v^m(\hat{S}) - p(t, \hat{S})$ and $v^m(S) - p(t+1, S) = v^m(\hat{S}) - p(t+1, \hat{S})$ for any $S, \hat{S} \in D^m(t) \cap D^m(t+1)$. These equalities imply that if bidders are truthful, then $p(t+1, S) - p(t, S) = p(t+1, \hat{S}) - p(t, \hat{S})$ for any $S, \hat{S} \in D^m(t) \cap D^m(t+1)$. Hence, in this case maximization in (6.12) is unnecessary.

We denote by Q^m the sum of $q^m(t)$ in the course of the auction, i.e., $Q^m = \sum_{t=1}^{T-1} q^m(t)$, where T denotes the step at which the auction terminates. We next show that Q^m reveals useful information about valuations of bidders, when they bid truthfully.⁵ A proof of this

⁵A similar result (with a slightly different definition for the quantity Q^m) that requires having piecewise

Lemma can be found in Section 6.8.

Lemma 6.4.6. *Assume that bidder m bids truthfully in the one-stage auction game, and the auction terminates assigning bundle S^m to bidder m . Then, $Q^m = v^m(\mathcal{N}) - v^m(S^m) + p^m(S^m)$, where $p^m(S^m)$ stands for the price associated with this bundle at the end of the auction.*

We next introduce the serial auction game. Let $\mathcal{M} = \{1, \dots, M\}$ denote the set of all bidders. The serial auction game involves:

- At stage zero, running a one-stage auction with bidders \mathcal{M} (referred to as auction 0),
- At stage k running a one-stage auction with bidders $\mathcal{M} - \{k\}$ (referred to as auction k).

Let Q_k^m denote the Q^m value associated with bidder m in auction k . Denote the assignment obtained at the end of stage k by $\{S_k^m\}$, and the corresponding prices for these bundles at the end of the auction by $\{p_k^m(S_k^m)\}$. If the serial auction game terminates, then

- Items are assigned according to the allocation that emerges at the end of stage 0, i.e. $\{S_0^m\}$.
- Final payment of each bidder m is equal to $\sum_{k \neq m} (Q_0^k - p_0^k(S_0^k)) - (Q_m^k - p_m^k(S_m^k))$.

Observe that the final allocation of items to bidders is determined at the end of stage 0. The subsequent stages are present to compute the final payments of the bidders.

A complete description of the game requires specifying payments of bidders for cases where the auction does not terminate. If auction l does not terminate ($l > 0$), then each bidder $m < l$ receives items as identified at the end of stage 0 ($\{S_0^m\}$), and has a final payment of $\sum_{k \neq m} (Q_0^k - p_0^k(S_0^k)) - (Q_m^k - p_m^k(S_m^k))$. Each bidder $m \geq l$, on the other hand, receives no items and has a payment equal to $\sum_{k \neq m} (Q_0^k - p_0^k(S_0^k))$. If auction 0 does not terminate, then no item is assigned, and the final payment of each bidder is equal to $\sup_T \sum_{k \neq m} \sum_{t=1}^T q^k(t)$.

We conclude this section by showing that it is an ex-post perfect equilibrium for bidders to bid truthfully in the serial auction game. Moreover, truthful bidding leads to an efficient allocation of items to bidders.

Theorem 6.4.1. *It is an ex-post perfect equilibrium for bidders to bid truthfully in the serial auction game. Moreover, the corresponding final allocation is efficient, and payments are the associated VCG payments.*

The proof of this theorem is provided in Section 6.8. Running a series of auctions may pose difficulties in practice. On the other hand, we demonstrate in Section 6.5 that when bidder-specific graphical pricing is used, the efficient outcome can be implemented by running a single iterative auction.

smooth and continuous price paths, and relies on computing a Stieltjes integral of prices is also present in Ausubel (2006).

Remark: Recall that the VCG payment of agent m acquiring bundle S^m of items is given by $\hat{\gamma}^m(S^m, \{v^k\}_{k \neq m}) \triangleq \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset} \sum_{k \neq m} v^k(Z^k) - \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset, Z^k \cap S^m = \emptyset} \sum_{k \neq m} v^k(Z^k)$. It follows from Lemma 6.4.6 that when agents $k \neq m$ bid truthfully in the one stage auction game, the auctioneer can compute the quantity $Q^k - p^k(S^m) = v^k(\mathcal{N}) - v^k(S^k)$ for all $k \neq m$. Since when bidders $k \neq m$ bid truthfully $\{S^k\}_{k \neq m}$ is an efficient allocation of items $\mathcal{N} - S^m$ to bidders $k \neq m$ (see Lemma 6.4.5), it follows that $\sum_{k \neq m} Q^k - p^k(S^m) = \sum_{k \neq m} v^k(\mathcal{N}) - v^k(S^k) = \sum_{k \neq m} v^k(\mathcal{N}) - \max_{\{Z^k\} | Z^k \cap Z^l = \emptyset, Z^k \cap S^m = \emptyset} \sum_{k \neq m} v^k(Z^k)$. This implies that $\hat{\gamma}^m(S^m, \{v^k\}_{k \neq m}) = \kappa^m + \sum_{k \neq m} (Q^k - p^k(S^m))$, where κ^m is a function of true valuations of bidders $k \neq m$. Thus, by running a single iterative auction, the auctioneer can learn the VCG payment of each agent (whose opponents truthfully bid) up to an additive function of her opponents' valuations. It can be shown that if this quantity is charged to bidders, as their final payments, then bidders still have no incentive to deviate from the truthful bidding strategy. However, some bidders may have negative payoffs at the end of the auction, and this may deter them from participating in the auction. Thus, if we allow outcomes where some bidders have negative payoffs, it may be possible to implement the efficient outcome by using a single iterative auction that relies on anonymous item-pricing.

An alternative approach involves running a single iterative auction, but not immediately terminating the auction when a Walrasian equilibrium is identified. In particular, the auctioneer can continue running the auction until she acquires sufficient information to compute the VCG payments (or identifies the efficient allocation for sets of bidders \mathcal{M} and $\mathcal{M} - \{k\}$ for all $k \in \mathcal{M}$). The iterative auction format we present in the next section makes use of this idea to implement the efficient auction by running a single iterative auction.

6.5 Iterative Auctions for General Graphical Valuations

In this section, we focus on solutions of LP5/D5 using iterative algorithms, and employ these to develop new iterative auction formats that rely on bidder-specific graphical pricing. The iterative auctions we develop in this section, have an ex-post perfect equilibrium that implements the efficient outcome for all graphical valuations. Importantly, unlike the auctions in Section 6.4, in these auctions, the auctioneer can compute the final payments that guarantee truthful bidding, by conducting only a single auction (as opposed to a series of auctions). Thus, our results imply that it suffices to run a single iterative auction which relies on bidder-specific graphical prices to implement the efficient outcome for graphical valuations.

An algorithm similar to that of Section 6.4 can be used for iteratively solving LP5 and finding the efficient allocation for graphical valuations. That is, we can start with a feasible dual solution, and check if there exists a primal feasible solution satisfying the complementary slackness condition with this dual solution. If such a solution exists, then we identify a primal-dual optimal pair. Otherwise, we can focus on a restricted primal problem

(similar to (6.10)), and the corresponding dual, and update the dual prices in accord with the solution of these optimization problems. Since, primal-dual algorithms converge to the optimal solution of the underlying optimization problem, and in LP5 there exists an optimal solution that coincides with the efficient allocation, the approach we outline above guarantees convergence to the efficient outcome.

Note that this approach requires solving the restricted primal/dual problems in order to identify the dual variable update directions. On the other hand, the dual space of LP5 (or the feasible set of D5) exhibits a special structure that allows for convergence to the optimal solution, by updating the dual variables using a simple update rule that does not require solving the restricted primal/dual problems. Before we explain this approach, we introduce some necessary definitions, and present important properties of the optimal solutions of D5.

We define a *restriction of a feasible solution* $(\{p_i^m, p_{ij}^m\}_{m \in \mathcal{M}}, \{\pi^m\}_{m \in \mathcal{M}}, \pi_s)$ of D5 to a set of bidders $\mathcal{M}_0 \subset \mathcal{M}$ as the tuple $(\{p_i^m, p_{ij}^m\}_{m \in \mathcal{M}_0}, \{\pi^m\}_{m \in \mathcal{M}_0}, \pi_s)$. Similarly, we refer to the tuple $\{p_i^m, p_{ij}^m, \pi^m\}_{m \in \mathcal{M}_0}$ as *the restriction of prices and bidder surpluses* to a set of bidders $\mathcal{M}_0 \subset \mathcal{M}$. Additionally, we say that *D5 is formulated with subset \mathcal{M}_0 of bidders*, if (i) it has p_i^m, p_{ij}^m, π^m variables only for $m \in \mathcal{M}_0$, (ii) the first constraint is imposed only for $m \in \mathcal{M}_0$, and (iii) the second constraint is present only for allocations μ , where bidders $m \notin \mathcal{M}_0$ do not receive any items, i.e., $\mu^m = \emptyset$.

We next show that there exists an optimal solution of D5, whose restriction to any subset of bidders gives prices and bidder surpluses that appear at an optimal solution of a formulation of D5, for this set of bidders.

Lemma 6.5.1. *(i) Assume that there are at least two bidders and one item, and the efficient allocation is unique. Then, the prices that are part of an optimal solution of D5 are not unique.*

(ii) There exists an optimal solution of D5, such that for any $\mathcal{M}_0 \subset \mathcal{M}$, the restriction of prices and bidder surpluses of this optimal solution to \mathcal{M}_0 agrees with prices and bidder surpluses at an optimal solution of a formulation of D5 with set of bidders \mathcal{M}_0 .

Proof. (i) Denote the unique efficient allocation by $\{\hat{S}^m\}$. Consider a solution of D5, where $p_i^m = w_i^m$, $p_{ij}^m = w_{ij}^m$, $\pi^m = 0$, for $m \in \mathcal{M}$, and π_s equals to the maximum welfare, i.e., $\pi_s = W^* \triangleq \max_{\{S^m\}} \sum_k v^k(S^k)$. It can be immediately checked that this solution is feasible in D5. Additionally, the corresponding objective value is equal to W^* . However, LP5 has a feasible solution, associated with the efficient allocation, which has the same objective value. This implies that the constructed dual feasible solution is also optimal.

Since the efficient allocation is unique, it follows that

$$\pi_s = \sum_m \sum_{i \in \hat{S}^m} p_i^m + \sum_{ij \in \hat{S}^m} p_{ij}^m > \sum_m \sum_{i \in S^m} p_i^m + \sum_{ij \in S^m} p_{ij}^m \quad (6.13)$$

for any other allocation $\{S^m\} \neq \{\hat{S}^m\}$. For some player k , we have $S^k - \hat{S}^k \neq \emptyset$. Let item j

belong to this set. For sufficiently small ϵ , consider the following solution of D5: $\bar{p}_j^k = p_j^k + \epsilon$, $\bar{p}_i^m = p_i^m$ for all $(m, i) \neq (k, j)$, $\bar{p}_{ir}^m = p_{ir}^m$ for all $(i, r) \in E$, $\bar{\pi}^m = \pi^m$ for all m , $\bar{\pi}^s = \pi^s$. Note that this solution is feasible since (weakly) increasing $\{p_i^m, p_{ij}^m\}$ does not violate the constraints involving π^m . Additionally, the constraint involving π^s is not violated since this constraint is strict except for allocation $\{\hat{S}^m\}$, $j \in S^m - \hat{S}^m$, and ϵ is sufficiently small. On the other hand, it can be seen that the new solution achieves the same objective value as the original solution, hence it is optimal, and the first part of the claim follows.

(ii) In order to prove the second part, consider the initial dual optimal solution we constructed for the first part, i.e., $p_i^m = w_i^m$, $p_{ij}^m = w_{ij}^m$, $\pi^s = \max_{\{S^m\}} \sum_k v^k(S^k)$, $\pi^m = 0$ for all players m , nodes i and edges $(i, j) \in E$. Observe that after restricting this solution to a set of bidders \mathcal{M}_0 , and replacing π^s with $\pi^s = \max_{\{S^m\} | S^m = \emptyset \text{ for } m \notin \mathcal{M}_0} \sum_k v^k(S^k)$, we obtain a solution for D5 formulated with bidders \mathcal{M}_0 . Feasibility of this solution can immediately be checked. Note that the objective value of D5 associated with this solution, is equal to the maximum welfare that can be raised with set of bidders \mathcal{M}_0 . On the other hand, when D5 is formulated with set of bidders \mathcal{M}_0 , the corresponding LP5 has a feasible solution associated with the efficient allocation to set of bidders \mathcal{M}_0 , and the objective value of this solution is equal to the maximum welfare ($\max_{\{S^m\} | S^m = \emptyset \text{ for } m \notin \mathcal{M}_0} \sum_k v^k(S^k)$). Thus, it follows that the solution we constructed is also optimal in a formulation of D5, with set of bidders \mathcal{M}_0 . We conclude that the restriction of prices and bidder surpluses of the optimal solution obtained in the first part to set \mathcal{M}_0 , agrees with an optimal solution to D5, formulated with set of bidders \mathcal{M}_0 , and the claim follows. \square

Since generically (i.e., except for possibly a measure zero set of node/edge weights) there is a unique efficient allocation, the above lemma suggests that D5 almost always has multiple optimal solutions. Moreover, it has an optimal solution whose prices and bidder surpluses also appear in an optimal solution of D5 formulated with fewer bidders.

Consider an optimal solution of D5, such that for all $m \in \mathcal{M}$, the prices and bidder surpluses of this solution agree with the prices and bidder surpluses at an optimal solution of D5, formulated with bidders $\mathcal{M} - \{m\}$. We refer to such a solution of D5 as a *special optimal solution*.⁶ Intuitively, the dual prices and bidder surpluses in this solution remain to be optimal, after removing at most one bidder from the formulation. Lemma 6.5.1 implies that special optimal solutions exist.

Let $(\{p_i^m\}, \{p_{ij}^m\}, \{\pi^m\}, \pi^s)$ be a special optimal solution of D5. By Theorem 5.4.1 it follows that LP5 has an associated integral optimal solution. Moreover, this solution

⁶Special optimal solutions are closely related to the universal competitive equilibrium (UCE) price concept of Mishra and Parkes (2007, 2009). In particular, UCE prices correspond to “competitive equilibrium” prices for sets of bidders \mathcal{M} and $\mathcal{M} - \{k\}$ (for every k). Moreover, they can be used for computing the VCG payments of agents. On the other hand, UCE prices associate a bidder-specific price with every bundle of items, and hence potentially consist of $M2^N$ distinct parameters. In contrast, the special optimal solutions we focus on here, associate a bidder-specific price with each node and edge of the underlying graph (hence consist of $O(MN^2)$ parameters). Moreover, when valuations are graphical, as we establish in this section, they are sufficient for the computation of VCG payments.

suggests an efficient allocation of items to bidders. We denote this efficient allocation by $\{S_0^m\}_{m \in \mathcal{M}}$, and note that the optimal solution of LP5 has $x^m(S_0^m) = 1$ for all m . Similarly, when D5 is formulated with a set of bidders $\mathcal{M} - \{k\}$, the corresponding LP5 has an integral optimal solution that identifies an efficient allocation for these bidders. For every $k \in \mathcal{M}$, denote this allocation by $\{S_k^m\}_{m \in \mathcal{M} - \{k\}}$, and note that the associated optimal solution of LP5 has $x^m(S_k^m) = 1$ for all $m \neq k$. Since $(\{p_i^m\}, \{p_{ij}^m\}, \{\pi^m\}, \pi^s)$ is a special optimal solution, then by complementary slackness in LP5 and D5, it follows that

$$\pi^m = v^m(S_0^m) - p^m(S_0^m) = v^m(S_k^m) - p^m(S_k^m),$$

where we use the shorthand notation $p^m(S) = \sum_{i \in S} p_i^m + \sum_{ij \in S} p_{ij}^m$. This suggests that at a special optimal solution we have:

$$v^m(S_k^m) - v^m(S_0^m) = p^m(S_k^m) - p^m(S_0^m). \quad (6.14)$$

Since by definition $\{S_k^m\}$ is an efficient allocation for a set of bidders $\mathcal{M} - \{k\}$, it follows that VCG payment (see Definition 6.3.3) for bidder k is given by $\sum_{m \neq k} v^m(S_k^m) - v^m(S_0^m) = \sum_{m \neq k} p^m(S_k^m) - p^m(S_0^m)$. This expression suggests that if a special optimal solution of D5 can be found, then this solution can be used to compute the VCG payments, in addition to the efficient outcome. Additionally, complementary slackness implies that at a special optimal solution the prices and the allocation $\{S_0^m\}$ constitute a pricing equilibrium with a bidder-specific graphical pricing rule.

We next propose an iterative auction format (see the table below) that implements the efficient allocation for all graphical valuations. The prices in this auction are updated so that when bidders reveal their demand truthfully, they converge to prices that are a part of a special optimal solution of D5. This allows for charging bidders final payments that are VCG payments, and guarantees that truthful demand revelation is an ex-post perfect equilibrium.

We make two assumptions before we state our auction format: (i) the valuation functions are integer-valued, (ii) there exists some integer \bar{w} such that $\bar{w} \geq w_{ij}^m, w_i^m$ for all m, i , and $(i, j) \in E$. Observe that the first assumption also implies that the node and edge weights are integers for all bidders. The second assumption simply suggests that an upper bound on node/edge weights is known.

This auction starts with high prices, at which bidders who bid truthfully do not demand any items. Intuitively, the auctioneer decreases the prices of items that are not demanded in isolation (i.e., $\{i\} \notin D^k$), until bidders start demanding them. Once a bidder demands such an item i , it is added to set Ψ^k . If two end points of an edge belong to this set (they are demanded at some point), and the bundle $\{i, j\}$ is not demanded, then the auctioneer decreases the price associated with this edge. If the termination condition in Step S2 holds, then the auction terminates with a final allocation, and payments that are only a function

An iterative auction for general graphical valuations.

S1: Set $p_i^m = p_{ij}^m = \bar{w}$ for all m, i and ij . Set $\Psi^m = \emptyset$ for all m .

S2: Ask every bidder m the sets she demands at the current prices, i.e., $D^m = \arg \max_S v^m(S) - \sum_{i \in S} p_i^m - \sum_{ij \in S} p_{ij}^m$.

If there exists an allocation $\{S_0^m\}$ such that

- $S_0^m \in D^m$
- $\sum_m p^m(S_0^m) \geq \sum_m p^m(S^m)$ for any other allocation $\{S^m\}$

and allocations $\{S_k^m\}_{m \neq k}$ for every bidder k such that

- $S_k^m \in D^m$
- $\sum_{m \neq k} p^m(S_k^m) \geq \sum_{m \neq k} p^m(S^m)$ for any other allocation $\{S^m\}$

then go to step S5. Otherwise go to step S3.

S3: For every bidder m , and item i , if $\{i\} \in D^m$, then do not update p_i^m , and set $\Psi^m := \Psi^m \cup \{i\}$. Otherwise decrement p_i^m by one.

For every bidder m , and edge $(i, j) \in E$, If $i, j \in \Psi^m$, and $\{i, j\} \notin D^m$, then decrement p_{ij}^m by one. Otherwise do not update p_{ij}^m .

S4: For every bidder m , if (i) $p_i^m < 0$ for some $i \in \mathcal{N}$, or (ii) $S \in D^m$ but $S' \subset S$ is such that $S' \notin D^m$, then set $p_i^m = p_{ij}^m = \bar{w}$ for all $i \in \mathcal{N}$, $(i, j) \in E$, and $\Psi^m = \emptyset$.

Go to Step S2.

S5: Terminate, by allocating items according to $\{S_0^m\}$, and assigning a final payment for each bidder k that is equal to $\sum_{m \neq k} p^m(S_k^m) - p^m(S_0^m)$.

of the prices that emerge in the last stage of the auction. The final payment of each bidder k is the difference between the revenue the auctioneer can raise at the final prices, if this bidder is not present ($\sum_{m \neq k} p^m(S_k^m)$) and the revenue the auctioneer can raise from the remaining bidders, when she is present ($\sum_{m \neq k} p^m(S_0^m)$). Finally, if a bidder m does not bid truthfully, then the conditions of Step S4 may hold, and the auctioneer “resets” the prices for bidder m , i.e., they are updated to the levels in Step S1. Note that if the auction does not terminate, the auctioneer does not allocate any items, and bidders do not make any payments. Hence, in this case we assume that bidders receive a payoff of zero.

An iterative solution of LP5/D5 using a primal-dual algorithm requires keeping a dual feasible solution at each stage, checking if there exists a primal feasible solutions that satisfies the complementary slackness conditions with the given dual feasible solution, and updating the dual variables to obtain a dual feasible solution with an improved objective. Additionally, as in Section 6.4, it is possible to implicitly define $\{\pi^m, \pi^s\}$ variables at each stage of this iterative algorithm, by setting $\pi^m = \max_S v^m(S) - \sum_{i \in S} p_i^m - \sum_{ij \in S} p_{ij}^m$, and $\pi^s = \max_{\mu \in \mathcal{X}} \sum_m \sum_{i \in \mu^m} p_i^m + \sum_{ij \in \mu^m} p_{ij}^m$. Note that this choice of dual variables guarantees

dual feasibility at each stage, and allows expressing dual variable updates, only in terms of the updates of the price variables.

The steps of the above iterative auction are similar to a primal-dual algorithm, when bidders truthfully reveal their demand. In particular, defining π^m , and π^s variables implicitly, as given above, dual feasibility in D5 is always guaranteed. Additionally, the first set of termination conditions in Step S2 (those involving S_0^m), correspond to checking if there exists a primal feasible solution satisfying complementary slackness conditions with the given dual feasible solution. More precisely, it can be seen that if these conditions hold, the primal feasible solution $x^m(S_0^m) = 1$ for all m , $\delta_{\{S_0^m\}} = 1$, and $x^m(S) = \delta_\mu = 0$ for remaining μ and (m, S) , satisfies complementary slackness conditions with the aforementioned dual feasible solution. Additionally, this solution is feasible in the primal, implying that the given primal and dual feasible solutions are optimal, and $\{S_0^m\}$ is efficient.

On the other hand, the iterative auction proposed above is different than a primal-dual algorithm in two aspects. First, even after an optimal solution is found the auction may not terminate. This can be seen by observing that if the first set of conditions in Step S2 hold, but the second set of conditions do not, the price updates continue. This feature of the auction guarantees that price updates terminate only when a special optimal solution is found (and possibly after an optimal solution is found). Secondly, unlike primal-dual algorithms, the dual updates need not strictly improve the objective. This can be seen by noting that even after an optimal solution is found, price updates continue. Moreover, price updates do not require explicitly formulating and solving restricted primal/dual problems.

Thus, we conclude that when bidders reveal their demand truthfully, the above iterative auction imitates a primal-dual algorithm. However, it is slightly different from a primal-dual algorithm, as it searches for a special optimal solution of D5 (as opposed to any optimal solution).

We next show that when bidders truthfully reveal their demand this iterative auction converges to a special outcome of D5. Additionally, when final allocation and payments are chosen as in step S5, it is an equilibrium for bidders to reveal their demand truthfully.

Theorem 6.5.1. *(i) Let $\{v^m\}$ denote the valuations of bidders. Assume that the above iterative auction terminates with prices $(\{p_i^m\}, \{p_{ij}^m\})$, and at the last step of the auction, bidders' demand reports are truthful. The solution of D5 obtained by setting $\pi^m = \max_S v^m(S) - \sum_{i \in S} p_i^m - \sum_{ij \in S} p_{ij}^m$, and $\pi^s = \max_{\mu \in \mathcal{X}} \sum_m \sum_{i \in \mu^m} p_i^m + \sum_{ij \in \mu^m} p_{ij}^m$, is a special optimal solution.*

(ii) In this auction, it is an ex-post perfect equilibrium for bidders to reveal their demand truthfully. Moreover, the corresponding final allocation is efficient, and payments are the associated VCG payments.

We conclude that the iterative auction format that is defined in this section guarantees that the efficient outcome emerges at an ex-post perfect equilibrium. Additionally, it does

so without running a series of auctions, and by relying on a simple pricing rule: bidder-specific graphical pricing. In particular, this auction terminates at a pricing equilibrium (with bidder-specific pricing rule), which also is a special optimal solution of D5. The allocation at this solution is efficient, and the prices allow for the computation of the VCG payments. These final payments ensure that bidders have no incentive to deviate from truthful bidding, at any stage of the iterative auction.

The key enabler of our result is the special structure of graphical valuation functions, and the existence of an LP formulation (LP5), which has a simple dual space, and can be used to find the efficient allocation for all graphical valuations. In the next section, we provide a generalization of our results and auction format to additively decomposable valuations, and discuss how the “complexity” of the auction format changes, as we consider more and more general valuation functions.

6.6 Generalization: An Alternative LP and Additively Decomposable Valuations

In the previous section we showed that D5 has some optimal solutions, which allow for computing the VCG payments. We designed an efficient iterative auction by ensuring convergence of prices to this special optimal solution. However, we also established that the optimal solution is not unique. Hence, the iterative auction we developed relied on convergence to the “right” optimal solution of D5.

In this section, we provide an alternative LP formulation of the efficient allocation problem, whose dual both suggests a bidder-specific graphical pricing rule, and reveals enough information at *all* of its optimal solutions to compute VCG payments. Iterative solutions of this LP formulation provide an alternative approach for designing iterative auctions systematically. Additionally, we show that it is possible to extend this approach to additively decomposable valuation functions introduced in Section 5.5.

Before presenting the alternative formulation, we introduce some new notation. We denote by χ_k all complete allocations where bidder k does not receive any items, i.e., if $\mu_k \in \chi_k$ then $\mu_k^k = \emptyset$. We use the variables $\{\delta_{\mu_k}^k\}_{\mu_k \in \chi_k}$ such that $\sum_{\mu_k \in \chi_k} \delta_{\mu_k}^k = 1$, $\delta_{\mu_k}^k \geq 0$, to denote a distribution over these allocations.

Using this notation, the primal LP (LP6) we focus on in this section is provided next:

$$\begin{aligned}
& \max \sum_m \sum_S x^m(S) v^m(S) \\
& \text{s.t.} \sum_S x^m(S) \leq M \quad \forall m \\
& \sum_{S|i \in S} x^m(S) \leq \sum_{\mu \in \chi | i \in \mu^m} \delta_\mu + \sum_{k \neq m} \sum_{\mu_k \in \chi_k | i \in \mu_k^m} \delta_{\mu_k}^k \quad \forall i, m \\
\text{LP6 :} \quad & \sum_{S|ij \in S} x^m(S) = \sum_{\mu \in \chi | ij \in \mu^m} \delta_\mu + \sum_{k \neq m} \sum_{\mu_k \in \chi_k | ij \in \mu_k^m} \delta_{\mu_k}^k \quad \forall ij \in E, m \\
& \sum_{\mu \in \chi} \delta_\mu = 1 \\
& \sum_{\mu_m \in \chi_m} \delta_{\mu_m}^m = 1 \quad \forall m \\
& \delta_\mu \geq 0, x^m(S) \geq 0.
\end{aligned}$$

In this optimization problem, we jointly solve for the efficient allocation for bidders in \mathcal{M} , and for bidders in $\mathcal{M} - \{k\}$ for all $k \in \mathcal{M}$. We could formulate separate optimization problems to compute these efficient allocations: (i) for the case involving all players in \mathcal{M} , we could use variables $(x^m(S), \delta_\mu)$ and formulate LP5, (ii) for bidders in $\mathcal{M} - \{k\}$ we could use variables $(x^m(S), \delta_{\mu_k}^k)$ and appropriately reformulate LP5 for bidders in $\mathcal{M} - \{k\}$. Instead, LP6 solves all of these optimization problems jointly, by aggregating their constraints and coupling them. For instance, as opposed to imposing $\sum_S x^m(S) \leq 1$ as in the M separate formulations of LP5 involving bidder m (mentioned above), LP6 imposes a single constraint $\sum_S x^m(S) \leq M$. Consequently, in LP6 we have a single constraint for each bidder-item or bidder-edge pair. This ensures that the corresponding dual optimization problem (D6), presented next, leads to a graphical pricing rule.

$$\begin{aligned}
& \min \pi^s + \sum_m \pi_m^s + M \sum_m \pi^m \\
& \text{s.t.} \pi^m \geq v^m(S) - \sum_{i \in S} p_i^m - \sum_{i,j \in S | ij \in E} p_{ij}^m \quad \forall S, m \\
\text{D6 :} \quad & \pi^s \geq \sum_m \left(\sum_{i \in \mu^m} p_i^m + \sum_{i \in \mu^m} p_{ij}^m \right) \quad \forall \mu \in \chi \\
& \pi_m^s \geq \sum_{k \neq m} \left(\sum_{i \in \mu_k^m} p_i^k + \sum_{i \in \mu_k^m} p_{ij}^k \right) \quad \forall \mu_m \in \chi_m \\
& p_i^m, \pi^m \geq 0.
\end{aligned}$$

Similar to our interpretation for π^s , in the dual problem, the variable π_m^s can be interpreted as the revenue of the auctioneer when bidder m does not receive any items. The

interpretation of the remaining variables/constraints is exactly the same as in D5.

We next show that for graphical valuations LP6 can always be used to jointly identify the efficient outcome in problem instances with bidders \mathcal{M} as well as $\mathcal{M} - \{k\}$. Additionally, the corresponding dual optimal solution of D6 can always be used to identify the VCG payments.

Theorem 6.6.1. *Assume that bidders have graphical valuations. Let $\{S^m\}_m$ and $\{S_k^m\}_{m \neq k}$ (for all $k \in \mathcal{M}$) denote the efficient allocation for set of bidders \mathcal{M} and $\mathcal{M} - \{k\}$ respectively.*

(i) *LP6 always has an optimal solution that is integral. In this solution, we have $\delta_{\{S^m\}} = \delta_{\{S_k^m\}}^k = 1$ for all $k \in \mathcal{M}$, and $\delta_\mu = \delta_{\mu_m}^m = 0$ for the remaining μ, μ_m, m ; and $x^m(S) = |\{k \in \mathcal{M} | S = S_k^m\}| + 1_{S=S^m}$ for all m, S , where $1_{S=S^m}$ is an indicator variable that is equal to 1 if $S = S^m$, and 0 otherwise.*

(ii) *At any dual optimal solution of D5, for any bidders m, k , we have*

$$\begin{aligned} \pi^m &= v^m(S^m) - \sum_{i \in S^m} p_i^m - \sum_{ij \in S^m} p_{ij}^m = v^m(S_k^m) - \sum_{i \in S_k^m} p_i^m - \sum_{ij \in S_k^m} p_{ij}^m, \\ \pi^s &= \sum_m \left(\sum_{i \in S^m} p_i^m + \sum_{ij \in S^m} p_{ij}^m \right), \\ \pi_k^s &= \sum_{m \neq k} \left(\sum_{i \in S_k^m} p_i^m + \sum_{ij \in S_k^m} p_{ij}^m \right). \end{aligned} \tag{6.15}$$

(iii) *At any dual optimal solution, for any bidder m , the quantity*

$$\sum_{k \neq m} \left(\left(\sum_{i \in S_k^m} p_i^k + \sum_{ij \in S_k^m} p_{ij}^k \right) - \left(\sum_{i \in S^k} p_i^k + \sum_{ij \in S^k} p_{ij}^k \right) \right) \tag{6.16}$$

is equal to the VCG payment of bidder m , for acquiring bundle S^m of items.

Proof. (i) Denote by W^* and W_k^* the total value generated by the allocations $\{S^m\}$ and $\{S_k^m\}$ respectively. It can be immediately checked that the solution specified in the theorem statement is feasible, and the associated objective value is equal to $W^* + \sum_k W_k^*$. This suggests that the optimal solution of LP6 is lower bounded by $W^* + \sum_k W_k^*$.

Consider the dual solution $p_i^m = w_i^m, p_{ij}^m = w_{ij}^m, \pi^m = 0$, for all $m, i \in \mathcal{N}, (i, j) \in E$, and

$$\begin{aligned} \pi^s &= \max_{\mu \in \mathcal{X}} \sum_m \left(\sum_{i \in \mu^m} p_i^m + \sum_{ij \in \mu^m} p_{ij}^m \right) \\ \pi_m^s &= \max_{\mu_k \in \mathcal{X}_k} \sum_{m \neq k} \left(\sum_{i \in \mu_k^m} p_i^m + \sum_{ij \in \mu_k^m} p_{ij}^m \right). \end{aligned}$$

It follows from D6 that this solution is feasible. Moreover, by construction the corresponding dual objective value is equal to $W^* + \sum_k W_k^*$. Thus, the optimal objective of the dual problem is upper bounded by $W^* + \sum_k W_k^*$.

By strong duality it follows that the primal and dual feasible solutions we construct above are optimal for LP6 and D6 respectively. Hence, the first part of the theorem follows.

(ii) Consider the optimal solution to LP6 given in part (i) of the theorem. Complementary slackness suggests that any dual optimal solution satisfies the conditions in the second part of the theorem. Hence, the claim immediately follows.

(iii) It follows from part (ii) that $\left(\sum_{i \in S_m^k} p_i^k + \sum_{ij \in S_m^k} p_{ij}^k\right) = v^k(S_m^k) - \pi^k$, and similarly $\left(\sum_{i \in S^k} p_i^k + \sum_{ij \in S^k} p_{ij}^k\right) = v^k(S^k) - \pi^k$. Thus, the quantity in (6.16) is equivalent to

$$\sum_{k \neq m} (v^k(S_m^k) - \pi^k) - (v^k(S^k) - \pi^k) = \sum_{k \neq m} (v^k(S_m^k) - v^k(S^k)).$$

Since allocations $\{S^m\}$ and $\{S_k^m\}$ are efficient (for set of bidders \mathcal{M} and $\mathcal{M} - \{k\}$ respectively), the quantity in the right hand side is equal to the VCG payment of bidder m , and the claim follows. \square

This theorem suggests that iterative solutions of LP6/D6 via primal-dual algorithms can be used to identify the efficient outcome, and VCG payments at the same time. Moreover, solution of these LPs using primal-dual algorithms lead to natural iterative auction formats. In these auctions, the auctioneer sets bidder-specific graphical prices, and the bidders report the set of items that they demand, exactly as in Section 6.5. Given demand reports the auctioneer can solve restricted primal/dual problems associated with LP6/D6 to find a dual update direction (assuming π^m, π^s, π_m^s variables are implicitly defined, analogously to our approach in Sections 6.4, 6.5). When the complementary slackness conditions associated with these LPs hold (or restricted problems have objective value zero), the auctioneer terminates the auction by assigning the allocation $\{S^m\}$ (as defined in Theorem 6.6.1(i)) and final payments specified in Theorem 6.6.1(iii). Since by collecting demand reports, and updating the prices as explained above, the auctioneer essentially runs a primal-dual algorithm, it can be seen that after any history H_t if bidders truthfully reveal their demand, this auction terminates at a pricing equilibrium with the pricing rule suggested by D6, and an efficient outcome is identified. Moreover, it can be easily checked that if the auction terminates assigning bidder m some set of items \hat{S}^m , and if bidders $k \neq m$ bid truthfully, the payments suggested in Theorem 6.6.1(iii) will be equal to $\hat{\gamma}(\hat{S}^m, \{v^k\}_{k \neq m})$. Thus, it can be shown that truthful bidding strategy satisfies conditions of Theorem 6.3.1, and hence in this auction it is an ex-post perfect equilibrium for bidders to reveal their demand truthfully. Since the approach outlined here for developing an iterative auction using a primal-dual algorithm, closely follows the approaches in Sections 6.4 and 6.5, the details are omitted.

We conclude this section by providing a generalization of LP6 and D6 to additively

decomposable valuations (denoted by *LP6G* and *D6G*). In particular, for a given collection $\mathcal{B} = \mathcal{B}_0 \cup \mathcal{B}_+ \cup \mathcal{B}_-$ of sets (see Section 5.5 for a definition of collections $\mathcal{B}_0, \mathcal{B}_+, \mathcal{B}_-$), we consider the following primal-dual LP pair:

$$\begin{aligned}
& \max \quad \sum_m \sum_S x^m(S) v^m(S) \\
& \text{s.t.} \quad \sum_S x^m(S) \leq M \quad \forall m \\
& \quad \sum_{S|S'=S \cap B} x^m(S) \leq \sum_{\mu|\mu^m \cap B=S'} \delta_\mu + \sum_{k \neq m} \sum_{\mu_k \in \chi_k | \mu_k^m \cap B=S'} \delta_{\mu_k}^k \quad \forall m, S' \subset B, B \in \mathcal{B}_+ \\
& \quad \sum_{S|S'=S \cap B} x^m(S) \geq \sum_{\mu|\mu^m \cap B=S'} \delta_\mu + \sum_{k \neq m} \sum_{\mu_k \in \chi_k | \mu_k^m \cap B=S'} \delta_{\mu_k}^k \quad \forall m, S' \subset B, B \in \mathcal{B}_- \\
\text{LP6G :} \quad & \sum_{S|S'=S \cap B} x^m(S) = \sum_{\mu|\mu^m \cap B=S'} \delta_\mu + \sum_{k \neq m} \sum_{\mu_k \in \chi_k | \mu_k^m \cap B=S'} \delta_{\mu_k}^k \quad \forall m, S' \subset B, B \in \mathcal{B}_0 \\
& \sum_{\mu \in \chi} \delta_\mu \leq 1 \\
& \sum_{\mu_k \in \chi_k} \delta_{\mu_k}^k \leq 1 \quad \forall k \\
& x^m(S), \delta_\mu, \delta_{\mu_k}^k \geq 0 \quad \forall m, S, \mu, \mu_k.
\end{aligned}$$

$$\begin{aligned}
& \min \quad \pi^s + \sum_m \pi_s^m + M \sum_m \pi^m \\
& \text{s.t.} \quad \pi^m \geq v^m(S) - \sum_B p_B^m(S \cap B) \quad \forall m, S \\
& \quad \pi^s \geq \sum_m \sum_B p_B^m(\mu^m \cap B) \quad \forall \mu \in \chi \\
& \quad \pi_s^m \geq \sum_{k \neq m} \sum_B p_B^k(\mu_k^m \cap B) \quad \forall m, \mu_m \in \chi_m \\
& \quad \pi^m \geq 0 \quad \forall m, \\
& \quad p_B^m(\cdot) \geq 0 \text{ for } B \in \mathcal{B}_+, \quad p_B^m(\cdot) \leq 0 \text{ for } B \in \mathcal{B}_-,
\end{aligned}$$

Observe that LP6G and D6G are immediate generalizations of LP6 and D6 respectively. In particular, the latter optimization formulations can be obtained by restricting attention to a formulation of LP6G and D6G, for a collection \mathcal{B} that contains only singletons, and pairs that correspond to edges.

Analogous to the results of Theorem 6.6.1, we show that this primal-dual LP pair can be used to find the efficient allocation, and the VCG payments for additively decomposable valuations with a given collection \mathcal{B} .

Theorem 6.6.2. *Assume that bidders have additively decomposable valuations with a collection \mathcal{B} . Let $\{S^m\}_m$ and $\{S_k^m\}_{m \neq k}$ (for all $k \in \mathcal{M}$) denote the efficient allocation for set*

of bidders \mathcal{M} and $\mathcal{M} - \{k\}$ respectively.

(i) LP6G always has an optimal solution that is integral. In this solution, we have $\delta_{\{S^m\}} = \delta_{\{S_k^m\}} = 1$ for all $k \in \mathcal{M}$, and $\delta_\mu = \delta_{\mu_m}^m = 0$ for the remaining μ, μ_m, m ; and $x^m(S) = |\{k \in \mathcal{M} | S = S_k^m\}| + 1_{S=S^m}$ for all m, S , where $1_{S=S^m}$ is an indicator variable that is equal to 1 if $S = S^m$, and 0 otherwise.

(ii) At any dual optimal solution of D6G, for any bidders m, k , we have

$$\begin{aligned}\pi^m &= v^m(S^m) - \sum_B p_B^m(S^m \cap B) = v^m(S_k^m) - \sum_B p_B^m(S_k^m \cap B), \\ \pi^s &= \sum_m \sum_B p_B^m(S^m \cap B), \\ \pi_k^s &= \sum_{m \neq k} \sum_B p_B^m(S_k^m \cap B).\end{aligned}$$

(iii) At any dual optimal solution, for any bidder m , the quantity

$$\sum_{k \neq m} \left(\sum_B p_B^k(S_m^k \cap B) - \sum_B p_B^k(S^k \cap B) \right) \quad (6.17)$$

is equal to the VCG payment of bidder m , for acquiring bundle S^m of items.

Proof. The proof of this theorem is identical to that of Theorem 6.6.1, and obtained following the same steps, after replacing p_i^m, p_{ij}^m by p_B^m , and constructing a feasible dual solution $p_B^m = w_B^m$. \square

This result suggests that using primal-dual algorithms with the aforementioned LP formulations, it is possible to jointly identify the efficient allocation, and VCG payments for additively decomposable valuations. Moreover, these algorithms suggest iterative auction formats that implement the efficient outcome for all graphical valuations. These auctions rely on using bidder-specific prices that decouple over the underlying collection of sets, i.e., $\{p_B^m\}$. Details are omitted, as this approach is an immediate generalization of our approach for graphical valuations, and closely follows our results in Sections 6.4 and 6.5.

6.7 Summary

In this chapter, we focused on iterative solutions of the LP formulations provided in Chapters 4 and 5, and showed that these can be used to design iterative auction formats where the auctioneer sets prices, bidders reveal their demand, and the auctioneer adjusts prices until a Walrasian/pricing equilibrium is found. When bidders truthfully reveal their demand, the final allocation of items to bidders identified by this process is efficient. We complemented

the iterative solutions of the LPs with appropriate payment schemes to guarantee truthfulness. Our results suggest that for sign-consistent tree valuations, the efficient allocation can be implemented (at an ex-post perfect equilibrium) by running a sequence of auctions that rely on the anonymous item pricing rule. On the other hand, if the auctioneer has the flexibility to use the bidder-specific graphical pricing rule, she can implement the efficient outcome (again at an ex-post perfect equilibrium) (i) for all graphical valuations, (ii) by running a single auction. In order for auctions that rely on bidder-specific graphical pricing rule to implement the efficient outcome, the auctioneer updates the prices in a way that leads convergence to a “special” dual optimal solution of the associated linear program. Motivated by this observation, we provided an alternative LP formulation of the efficient allocation problem for general graphical valuations. This formulation still suggests employing bidder-specific graphical prices, and its solution simultaneously reveals the efficient outcome (i) for all bidders, as well as (ii) for all bidders but one. Additionally, any dual optimal solution of this LP formulation can be used to compute the VCG payments. Thus, this LP can be used together with any primal-dual algorithm to develop iterative auctions that implement the efficient outcome. Moreover, this formulation generalizes to settings with additively decomposable valuations, providing a framework for developing iterative auction formats that guarantee efficiency beyond graphical valuations.

The results of this part of the thesis imply that when valuation functions of bidders exhibit some structure (such as the additively decomposable structure, or the graphical structure), it is possible to develop efficient iterative auction formats that rely on pricing rules which have a similar structure. Thus, the “complexity” of the pricing rule need not exceed the “complexity” of the valuation functions for iterative auction design. We close this part of the thesis by emphasizing that in practice it may be possible to develop iterative auction formats that rely on simple pricing rules, by first identifying the structure in valuations of bidders, and then following the framework provided here to exploit this special structure.

6.8 Appendix: Additional Proofs

Proof of Lemma 6.4.1. If $\{x^m(S)\}$ is a solution to (6.9), then it satisfies complementary slackness conditions, and this solution is optimal in LP1 and the associated dual solution (p, π) is optimal in D1. On the other hand, Corollary 4.3.1 suggests that there exists an optimal solution of LP1 $\{\bar{x}^m(S)\}$ that is integral. Observe that if $\sum_S \bar{x}^m(S) = 0$ for some m , by setting $x^m(\emptyset) = 1$, another integral feasible solution of LP1 with the same objective value can be obtained. Thus, without loss of generality we can assume that the optimal integral solution $\{\bar{x}^m(S)\}$ is such that $\sum_S \bar{x}^m(S) = 1$ for all m . Moreover, Assumption 4.2.1 implies that if there exists an optimal solution of LP1 that is integral, another integral optimal solution such that $\sum_m \sum_{S|i \in S} \bar{x}^m(S) = 1$ for all i can be obtained. This can be

seen by noting that given an optimal integral solution $\{\bar{x}^m(S)\}$ such that $\bar{x}^m(S^m) = 1$, and $\sum_m \sum_{S|i \in S} \bar{x}^m(S) = 0$ for node i , another feasible solution with weakly larger objective value can be obtained by setting $\bar{x}^m(S^m \cup \{i\}) = 1$, and $\bar{x}^m(S^m) = 0$, and keeping the remaining components of the original solution intact.

These observations imply that an optimal solution of LP1 that is integral and that satisfies $\sum_m \sum_{S|i \in S} \bar{x}^m(S) = 1$ for all i , and $\sum_S \bar{x}^m(S) = 1$ for all m can be obtained. Denote this solution by $\{\hat{x}^m(S)\}$. Since, (p, π) is optimal in D1, it follows that it should satisfy complementary slackness conditions with $\hat{x}^m(S)$. These conditions are identical to (6.9), and hence the claim follows. \square

Proof of Lemma 6.4.2. First, we show that if (p, π) is feasible in D1, then for small enough ϵ , so is $(p, \pi) + \epsilon(\bar{p}, \bar{\pi})$. To see this first observe that if $p_i > 0$, then for sufficiently small ϵ , we have $p_i + \epsilon\bar{p}_i \geq 0$. On the other hand, if $p_i = 0$, then feasibility of $(\bar{p}, \bar{\pi})$ in (6.11) implies that $\bar{p}_i \geq 0$, and hence $p_i + \epsilon\bar{p}_i \geq 0$. Similarly, observe that for all m, S such that $\pi^m = v^m(S) - \sum_{i \in S} p_i$, we have $\bar{\pi}^m + \sum_{i \in S} \bar{p}_i \geq 0$. This implies that $\pi^m + \epsilon\bar{\pi}^m \geq v^m(S) - \sum_{i \in S} (p_i + \epsilon\bar{p}_i)$. On the other hand, if $\pi^m > v^m(S) + \sum_{i \in S} p_i$, then for sufficiently small ϵ , we have $\pi^m + \epsilon\bar{\pi}^m \geq v^m(S) + \sum_{i \in S} (p_i + \epsilon\bar{p}_i)$. Thus we conclude that for all m and S this inequality holds. Moreover, for $S = \emptyset$, this inequality implies that $\pi^m + \epsilon\bar{\pi}^m \geq 0$. These observations imply that $(p, \pi) + \epsilon(\bar{p}, \bar{\pi})$ is a feasible solution of D1.

Second, we note that since (6.9) has no solution, the problem (6.10) and its dual (6.11) have positive optimal values, and consequently we have $-\sum_m \bar{\pi}^m - \sum_i \bar{p}_i > 0$ and $\sum_m \bar{\pi}^m + \sum_i \bar{p}_i < 0$. Thus, it follows that $(\bar{p}, \bar{\pi})$ is a valid dual update direction, and to complete the proof of the lemma, it suffices to establish that (i), (ii), and (iii) hold.

Consider an optimal solution of (6.10). If i is overdemanded, then at this solution we have $h_i > 0$. Using complementary slackness conditions in (6.10) suggests that at the corresponding dual optimal solution we have $\bar{p}_i = 1 > 0$. Similarly, if i is underdemanded, then in the optimal solution of (6.10), we have $\gamma_i > 0$. In this case the complementary slackness conditions imply that $\bar{p}_i = -1 < 0$. Hence, we conclude that the claims (i) and (ii) both hold, for some dual update that can be obtained by a solution of (6.10) and (6.11).

In order to prove (iii), assume that $\pi^m = \max_S v^m(S) - \sum_{i \in S} p_i$. Observe that in the optimal solution of (6.11), for every player m , and some set S^* , the first constraint is active (i.e., met with equality). This is because, otherwise $\bar{\pi}^m$ can be decreased to obtain a solution with better objective value. Thus, m, S^* satisfy $\pi^m = v^m(S^*) - \sum_{i \in S^*} p_i$, and $\bar{\pi}^m + \sum_{i \in S^*} \bar{p}_i = 0$. These imply that $\pi^m + \epsilon\bar{\pi}^m = v^m(S^*) - \sum_{i \in S^*} (p_i + \epsilon\bar{p}_i)$. \square

Proof of Lemma 6.4.3. (i) Since, $(\bar{p}, \bar{\pi})$ is an improvement direction by construction, it suffices to construct $\phi(p, \bar{p}) > 0$ such that $(p, \pi) + \phi(p, \bar{p})(\bar{p}, \bar{\pi})$ is feasible to establish the first part of the result.

Observe that if for some m and S , $\pi^m = v^m(S) - \sum_{i \in S} p_i$, then by dual feasibility in (6.11) we have $\bar{\pi}^m \geq \sum_{i \in S} \bar{p}_i$, and hence $\pi^m + \theta\bar{\pi}^m \geq v^m(S) - \sum_{i \in S} (p_i + \theta\bar{p}_i)$ for any

$\theta > 0$. Thus, feasibility constraint trivially follows for demanded sets $S^m \triangleq \{S \mid \pi^m = v^m(S) - \sum_{i \in S} p_i\}$, for any choice of θ .

Let $S \in S^m$ and $S' \notin S^m$ respectively be sets that are demanded and not demanded by player m at prices p i.e., $\pi^m = v^m(S) - \sum_{i \in S} p_i > v^m(S') - \sum_{i \in S'} p_i$. Observe that

$$\sum_{i \in S'} p_i - \sum_{i \in S} p_i > v^m(S') - v^m(S), \quad (6.18)$$

where the quantity in the right hand side is an integer, since valuations are integer-valued. Let

$$\Delta_{S,S'} \triangleq \begin{cases} 1 & \text{if } \sum_{i \in S'} p_i - \sum_{i \in S} p_i \text{ is an integer} \\ (\sum_{i \in S'} p_i - \sum_{i \in S} p_i) - \lfloor \sum_{i \in S'} p_i - \sum_{i \in S} p_i \rfloor & \text{otherwise,} \end{cases} \quad (6.19)$$

where $\lfloor z \rfloor$ designates the largest integer that is weakly smaller than z . Observe that by construction $\Delta_{S,S'} > 0$. Additionally, the definition of $\Delta_{S,S'}$ guarantees that $\sum_{i \in S'} p_i - \sum_{i \in S} p_i - \Delta_{S,S'} \geq v^m(S') - v^m(S)$. Rearranging terms, we obtain

$$v^m(S) - \sum_{i \in S} p_i \geq v^m(S') - \sum_{i \in S'} p_i + \Delta_{S,S'}. \quad (6.20)$$

Assume that θ is chosen so that $\theta(\sum_{i \in S'} \bar{p}_i - \sum_{i \in S} \bar{p}_i) \geq -\Delta_{S,S'}$. Then

$$\begin{aligned} \pi^m + \theta \bar{\pi}^m &\geq v^m(S) - \sum_{i \in S} (p_i + \theta \bar{p}_i) \\ &\geq v^m(S') - \sum_{i \in S'} p_i + \Delta_{S,S'} - \theta \sum_{i \in S} \bar{p}_i \\ &\geq v^m(S') - \sum_{i \in S'} p_i - \theta \sum_{i \in S'} \bar{p}_i = v^m(S') - \sum_{i \in S'} (p_i + \theta \bar{p}_i) \end{aligned} \quad (6.21)$$

where the second line follows from (6.20), and the last line follows from the assumption on θ . Note that this inequality suggests that the new dual solution satisfies the constraint $\pi^m + \theta \bar{\pi}^m \geq v^m(S') - \sum_{i \in S'} (p_i + \theta \bar{p}_i)$.

Thus, defining $\phi(p, \bar{p}) \triangleq \max \theta$ subject to

- $\theta(\sum_{i \in S'} \bar{p}_i - \sum_{i \in S} \bar{p}_i) \geq -\Delta_{S,S'}$ for all m and set S demanded by m , and S' not demanded by m ,
- $p_i + \theta \bar{p}_i \geq 0$,

it follows that the updated solution $(p, \pi) + \phi(p, \bar{p})(\bar{p}, \bar{\pi})$ remains feasible in D1. Note that by (6.19), we have $\Delta_{S,S'} > 0$. Additionally, by dual feasibility in (6.11) we obtain $\bar{p}_i = 0$ when $p_i = 0$. These imply $\phi(p, \bar{p}) > 0$, and the first part of the claim follows.

(ii) Assume that $\phi(p, \bar{p}) \neq \theta^*$, and $(\hat{p}, \hat{\pi}) = (p, \pi) + \phi(p, \bar{p})(\bar{p}, \bar{\pi})$. Since θ^* the largest real that guarantees feasibility of $(p, \pi) + \theta^*(\bar{p}, \bar{\pi})$ in D1, it follows that $\phi(p, \bar{p}) < \theta^*$.

By definition of θ^* it follows that if $p_i > 0$, then $\hat{p}_i > 0$, and if $\pi^m > v^m(S) - \sum_{i \in S} p_i$, then $\hat{\pi}^m > v^m(S) - \sum_{i \in S} \hat{p}_i$, as well. These observations suggest that $(\bar{p}, \bar{\pi})$ is feasible in a formulation of (6.11) associated with $(\hat{p}, \hat{\pi})$. Additionally they imply that constraints of (6.10) that are different under $(\hat{p}, \hat{\pi})$ (when compared to (p, π)) are associated with (a) nodes i , where $p_i = 0$, and $\hat{p}_i > 0$, or (b) bidder m and bundle S , where $\pi^m = v^m(S) - \sum_{i \in S} p_i$ but $\hat{\pi}^m > v^m(S) - \sum_{i \in S} \hat{p}_i$. In case (a), for node i we replace the constraint $\sum_m \sum_{S|i \in S} x^m(S) - h_i \leq 1$ (associated with solution (p, π)) with $\sum_m \sum_{S|i \in S} x^m(S) + \gamma_i - h_i = 1$ (associated with solution $(\hat{p}, \hat{\pi})$). In case (b) we include a new constraint $x^m(S) = 0$ for the solution $(\hat{p}, \hat{\pi})$. We next show that these changes in the constraint set do not change the optimal objective of (6.10).

Observe that in case (a), it should be the case that $\bar{p}_i > 0$. Since $(\bar{p}, \bar{\pi})$ is an optimal solution of (6.11), complementary slackness conditions imply that at optimal solutions of (6.10) associated with (p, π) we have

$$\sum_m \sum_{S|i \in S} x^m(S) - h_i = 1. \quad (6.22)$$

This implies that the optimal solution of (6.10) associated with the original dual solution (p, π) satisfies the new constraint associated with node i when we reformulate (6.10) for $(\hat{p}, \hat{\pi})$, namely $\sum_m \sum_{S|i \in S} x^m(S) + \gamma_i - h_i = 1$.

Similarly, in case (b) it should be the case that $\bar{\pi}^m > \sum_{i \in S} \bar{p}_i$ at an optimal solution of (6.11) associated with (p, π) . Complementary slackness implies that $x^m(S) = 0$ at the corresponding optimal solution of (6.10). This implies that the new constraint in (6.10) in case (b), namely $x^m(S) = 0$, is trivially satisfied by the optimal solution of (6.10) associated with the original dual solution (p, π) .

These observations imply that optimal solution of (6.10) associated with the original dual solution (p, π) , remains feasible (after complementing it with $\gamma_i = 0$ for i associated with case (a)) in the same problem reformulated according to $(\hat{p}, \hat{\pi})$. On the other hand, it follows that the new constraints added to (6.10) suggest that the optimal value of the new problem should be weakly larger than that of the original one. This is because the addition of constraint $x^m(S) = 0$ for some m and S (in case (b)) can only make the feasible set smaller, whereas the constraint $\sum_m \sum_{S|i \in S} x^m(S) + \gamma_i - h_i = 1$ replacing $\sum_m \sum_{S|i \in S} x^m(S) - h_i \leq 1$ includes an additional penalty term γ_i . Thus, the optimal solution of (6.10) associated with the original dual solution (p, π) is not only feasible but also optimal in the new problem associated with $(\hat{p}, \hat{\pi})$. This implies that $(\bar{p}, \bar{\pi})$, which is feasible in the dual problem (6.11) associated with $(\hat{p}, \hat{\pi})$ is optimal in the same problem. Hence, the second part of the claim follows.

(iii) Consider an initial dual feasible solution (p, π) , and the associated dual optimal

solution $(\bar{p}, \bar{\pi})$ of (6.11). Denote by $(p(r), \pi(r))$ the dual feasible solution obtained after r dual variable updates, starting from (p, π) . Observe that part (ii) suggests that at each stage the dual updates take the form $(p(r+1), \pi(r+1)) = \phi(p(r), \bar{p})(\bar{p}, \bar{\pi}) + (p(r), \pi(r))$, i.e., the dual update direction $(\bar{p}, \bar{\pi})$ is used at each stage.

Observe that the construction of $\phi(p(r), \pi(r))$ (presented in part (i)) suggests that after the dual update at stage r , either $p_i(r+1) = p_i(r) + \phi(p(r), \bar{p})\bar{p}_i = 0$ (while $p_i > 0$) for some i , or $\phi(p(r), \bar{p})(\sum_{i \in S'} \bar{p}_i - \sum_{i \in S} \bar{p}_i) = -\Delta_{S,S'}$ for some demanded set S , and undemanded set S' . If it is the former, it follows from the definition of θ^* and dual feasibility of $(p(r), \pi(r))$ at each stage that $\sum_{l=1}^r \phi(p(l), \bar{p}) = \theta^*$, and the claim immediately follows. Assume that it is the latter.

Observe that the definition of $\Delta_{S,S'}$ implies that in the latter case we have $\sum_{i \in S} p_i(r+1) - \sum_{i \in S'} \hat{p}_i(r+1)$ equal to an integer after the dual update. Thus, after the update $\Delta_{S,S'}$ becomes equal to 1 (see (6.19)). Since there are at most 2^N choices for S , and 2^N choices for S' , it follows that after at most 2^{2N} iterations the latter case holds for the same S, S' . These observations imply that after at most 2^{2N} iterations, $\kappa(\sum_{i \in S'} \bar{p}_i - \sum_{i \in S} \bar{p}_i) \leq -1$, where $\kappa = \sum_{l=1}^{2^{2N}} \phi(p(l), \bar{p})$. This implies that after finitely many steps for some S, S' the inequality (6.18) holds with equality. Let L denote the first iteration where this inequality holds with equality for some set S' that is not initially demanded. On the other hand, this suggests that S' (that is not demanded at the original solution) starts being demanded at stage L . Note that by definition θ^* is the smallest real number where at dual solution $(p, \pi) + \theta^*(\bar{p}, \bar{\pi})$, a new set starts to be demanded or $p_i + \theta^*\bar{p}_i = 0$ (while $p_i > 0$). Since $p_i(r) + \theta^*\bar{p}_i(r) > 0$ for all r , in our construction, and L is the first stage a new set starts to be demanded, it follows that $\theta^* = \sum_{l=1}^L \phi(p(l), \bar{p})$.

This implies that the solution reached after finitely many updates (using step size $\phi(p(r), \bar{p})$ at stage r) is $(p, \pi) + \theta^*(\bar{p}, \bar{\pi})$. Thus the claim follows. \square

Proof of Lemma 6.4.5. (i) If the one-stage auction game terminates with an allocation $\{S^m\}$, then the termination conditions in Step S3 imply that $v^m(S^m) - \sum_{i \in S^m} p_i \geq v^m(S) - \sum_{i \in S} p_i$ for $m \in \mathcal{M}_0$. Let $\{\hat{S}^m\}_{m \in \mathcal{M}_0}$ denote another allocation of items in $\cup_{m \in \mathcal{M}_0} S^m$ to bidders in \mathcal{M}_0 . Observe that

$$\sum_{m \in \mathcal{M}_0} \left(v^m(S^m) - \sum_{i \in S^m} p_i \right) \geq \sum_{m \in \mathcal{M}_0} \left(v^m(\hat{S}^m) - \sum_{i \in \hat{S}^m} p_i \right),$$

or equivalently

$$\sum_{m \in \mathcal{M}_0} \left(v^m(S^m) - v^m(\hat{S}^m) \right) \geq \sum_{i \in \cup_{m \in \mathcal{M}_0} S^m} p_i - \sum_{i \in \cup_{m \in \mathcal{M}_0} \hat{S}^m} p_i \geq 0, \quad (6.23)$$

where the inequality follows from $p_i \geq 0$ and $\cup_{m \in \mathcal{M}_0} \hat{S}^m \subset \cup_{m \in \mathcal{M}_0} S^m$. Since this is true for

any $\{\hat{S}^m\}_{m \in \mathcal{M}_0}$, the claim follows from (6.23).

(ii) The first part implies that if the one-stage auction game terminates, then the outcome is efficient. In order to prove the claim it suffices to prove that the process terminates when bidders reveal their demand truthfully, after history H_t , starting from time $t + 1$.

Observe that if bidders reveal their demand truthfully at time $t + 1$, given prices p , we can associate with this a dual feasible solution (p, π) of D1, where $\pi^m = v^m(S) - \sum_{i \in S} p_i$, for any set S that is demanded. On the other hand, by construction, one-stage auction updates the prices as suggested by Algorithm 1. This implies that from time $t + 1$ onwards, the demand sets and prices that are revealed in the one-stage auction game coincide with those of Algorithm 1. On the other hand, since it is a primal-dual algorithm, Algorithm 1 terminates with the efficient allocation starting from any dual feasible solution (Lemma 6.4.4). This implies that one stage auction game terminates with an efficient allocation of items to bidders. \square

Proof of Lemma 6.4.6. Recall that in the one-stage auction game price updates are identical to those in Algorithm 1, and the price update direction is obtained by solving (6.11). On the other hand, Lemma 6.4.2(iii) suggests that when bidder m reveals her demand truthfully the price updates are such that there is a bundle that is demanded by this bidder at time t and $t - 1$. This implies that if bidder m reveals her demand truthfully, then for all t , we have $q^m(t) = p(t + 1, S) - p(t, S)$ for $S \in D^m(t) \cap D^m(t + 1)$.

Let $\pi^m(t)$ denote the surplus bidder m has for a bundle she demands at time t , i.e., $\pi^m(t) = v^m(S) - \sum_{i \in S} p_i(t)$, for $S \in D^m(t)$. Using the shorthand notation $p^m(t, S) = \sum_{i \in S} p_i(t)$, this implies that $\pi^m(t) - \pi^m(t + 1) = p^m(t + 1, S) - p^m(t, S)$ for any $S \in D_t^m \cap D_{t+1}^m$. Thus, it follows that $Q^m = \pi^m(1) - \pi^m(T)$. Note that at step 1, prices are equal to zero, and hence $\mathcal{N} \in D^m(1)$. Consequently, substituting $\pi^m(1) = v^m(\mathcal{N})$, and $\pi^m(T) = v^m(S^m) - p^m(T, S^m)$, it follows that $Q^m = v^m(\mathcal{N}) - (v^m(S^m) - p^m(T, S^m))$, where $p^m(T, S^m)$ is the final price for bundle S^m , as the claim suggests. \square

Proof of Theorem 6.4.1. We prove the ex-post equilibrium result by establishing that after any history, no bidder can improve her payoff by deviating from the truthful bidding strategy, provided that her opponents bid truthfully.

Consider bidder m , and history H_t such that after time t , the auction is at stage $l > m$. From the definition of the serial auction game it follows that bidder m 's payoff is the same for any strategy (the allocation is already determined by auction 0, and her payment is determined by auction m). Hence, bidder m has no incentive to deviate from the truthful bidding strategy after H_t .

Assume instead that after history H_t , the auction is at stage l such that $0 < l < m$ (bidder m does not participate in auction m , so this case is excluded). Note that the allocation is determined at stage 0, and bidder m does not participate in auction m , where her final payment is determined. So, she can only impact her payoff in stages $0 < l < m$, by

using a strategy that prevents auction l from terminating. Let, u^m denote her payoff when she bids truthfully, and \hat{u}^m denote the payoff when she prevents termination of some stage $0 < l < m$. Since her opponents bid truthfully, it follows that (recall that $\{S_0^k\}$ and $\{S_m^k\}$ are the allocations that emerge at the end of auctions 0 and m respectively, when auction m terminates)

$$\begin{aligned}
u^m &= v^m(S_0^m) - \sum_{k \neq m} (Q_0^k - p_0^k(S_0^k)) - (Q_m^k - p_m^k(S_m^k)) \\
&= v^m(S_0^m) - \sum_{k \neq m} (Q_0^k - p_0^k(S_0^k)) - (v^k(\mathcal{N}) - v^m(S_m^k)) \\
&\geq v^m(S_0^m) - \sum_{k \neq m} (Q_0^k - p_0^k(S_0^k)) \geq - \sum_{k \neq m} (Q_0^k - p_0^k(S_0^k)) = \hat{u}^m,
\end{aligned} \tag{6.24}$$

where the second equality follows from Lemma 6.4.6 and the fact that bidders $k \neq m$ bid truthfully (in auction m), and the inequalities follow from the fact that $v^k(\mathcal{N}) - v^m(S_m^k) \geq 0$, and $v^m(S_0^m) \geq 0$. Thus, bidder m has no incentive to deviate from truthful bidding after any history H_t such that the auction is in stage $l \in \{1, \dots, m-1\}$ after H_t .

We complete the proof by showing that bidder m cannot deviate from truthful bidding and improve her payoff after some history H_t , such that after t the auction is still in stage 0. Observe that there are two cases to consider (a) she can use some strategy z^m , and auction 0 terminates with bundles $\{S_0^k\}$, or (b) auction 0 never terminates.

First consider case (a). From the definition of serial auction (and the fact that bidders bid truthfully for any history, after stage 0), it follows that bidder m 's payment (denoted by Γ^m) is equal to

$$\Gamma^m = \sum_{k \neq m} (Q_0^k - p_0^k(S_0^k)) - (Q_m^k - p_m^k(S_m^k)). \tag{6.25}$$

Denote by T_0 the time at which stage 0 terminates. By definition, for every bidder $k \neq m$ we have

$$Q_0^k = \sum_{r=1}^{T_0-1} q_0^k(r) = \sum_{r=1}^t q_0^k(r) + \sum_{r=t+1}^{T_0-1} q_0^k(r), \tag{6.26}$$

where q_0^k is defined according to (6.12) for auction 0. Let $S^k(r)$ denote a bundle demanded at step r of auction 0. Since bidders $k \neq m$ bid truthfully after time t , and in the auction prices are updated according to Algorithm 1, it follows from Lemma 6.4.2 (iii) that at all $r > t$ there exist some bundle $S^k(r) \in D^k(r) \cap D^k(r+1)$. Consequently, we obtain from (6.12) that $q_0^k(r) = p(r+1, S^k(r)) - p(r, S^k(r))$. Denoting the surplus of bidder k at step r of auction 0 by $\pi_0^k(r) = v^k(S^k(r)) - p(r, S^k(r))$, we obtain $q_0^k(r) = \pi_0^k(r) - \pi_0^k(r+1)$, since $S^k(r)$ is demanded both at time r and $r+1$. This implies that $\sum_{r=t+1}^{T_0-1} q_0^k(r) = \sum_{r=t+1}^{T_0-1} \pi_0^k(r) - \pi_0^k(r+1) = \pi_0^k(t+1) - \pi_0^k(T_0)$. Since $\pi^k(r)$ is the surplus of bidder k at r , and $S^k(r)$ is a demanded set, we can rewrite this expression as $\sum_{r=t+1}^{T_0-1} q_0^k(r) = (v^k(S^k(t+1)) - p^k(t+1, S^k(t+1))) - (v^k(S^k(T_0)) - p^k(T_0, S^k(T_0)))$. Together with (6.26),

this implies that

$$Q_0^k - p^k(T_0, S^k(T_0)) = \sum_{r=1}^t q_0^k(r) + v^k(S^k(t+1)) - p^k(t+1, S^k(t+1)) - v^k(S^k(T_0)). \quad (6.27)$$

Since bidders $k \neq m$ bid truthfully, it follows from Lemma 6.4.6 that $Q_m^k - p_m^k(S_m^k) = v^k(\mathcal{N}) - v^k(S_m^k)$, where $\{S_m^k\}$ denotes the allocation that emerges at the end of auction m . This equation, together with (6.25) and (6.27) implies that (note that by definition $p_0^k(S_0^k) = p^k(T_0, S^k(T_0))$),

$$\begin{aligned} \Gamma^m &= \sum_{k \neq m} \left(\sum_{r=1}^t q_0^k(r) + v^k(S^k(t+1)) - p^k(t+1, S^k(t+1)) - v^k(S^k(T_0)) \right) - \sum_{k \neq m} \left(v^k(\mathcal{N}) - v^k(S_m^k) \right) \\ &= \sum_{k \neq m} \left(v^k(S_m^k) - v^k(S^k(T_0)) - v^k(\mathcal{N}) \right) + \kappa_t^m, \end{aligned} \quad (6.28)$$

where $\kappa_t^m \triangleq \sum_{k \neq m} \left(\sum_{r=1}^t q_0^k(r) - p^k(t+1, S^k(t+1)) + v^k(S^k(t+1)) \right)$. Thus, in case (a) payoff of bidder m is equal to

$$u^m = v^m(S^m(T_0)) - \Gamma^m = \sum_k v^k(S^k(T_0)) - \sum_{k \neq m} \left(v^k(S_m^k) - v^k(\mathcal{N}) \right) - \kappa_t^m. \quad (6.29)$$

Observe that κ_t^m is a function of H_t , and hence is independent of bidder m 's strategy after step t . On the other hand, the quantity $\sum_{k \neq m} \left(v^k(S_m^k) - v^k(\mathcal{N}) \right)$ is only a function of the outcome of auction m , which is also independent of bidder m 's strategy. This implies that by choosing a different strategy, bidder m can modify only $\sum_k v^k(S^k(T_0))$ component of her payoff, provided that the auction terminates. On the other hand, Lemma 6.4.5 suggests that $\{S^k(T_0)\}_k$ is an efficient allocation of items to bidders $k \neq m$, when z^m is the truthful bidding strategy. Thus, any deviation from the truthful bidding strategy after stage H_t decreases bidder m 's payoff, provided that the auction terminates (i.e., case (a)).

On the other hand, if after H_t , bidder m uses a strategy that leads to nontermination in auction 0 (i.e., case (b)), then by definition of the serial auction game her payoff is

$$\hat{u}^m = - \sup_T \sum_{k \neq m} \sum_{r=1}^T q^k(r). \quad (6.30)$$

Thus, it follows that $\hat{u}^m \leq - \sum_{k \neq m} \sum_{r=1}^T q^k(r) \leq -\kappa_t^m + \sum_{k \neq m} v^k(S^k(t+1))$. Observe that if bidders bid truthfully after H_t , then $\{S_0^k\}$ is the efficient allocation. Thus, $\sum_{k \neq m} v^k(S_m^k) \leq \sum_{k \neq m} v^k(S_0^k)$. This observation and $v^k(S^k(t+1)) \leq v^k(\mathcal{N})$, imply that $\sum_{k \neq m} v^k(S^k(t+1)) \leq \sum_k v^k(S^k(T_0)) - \sum_{k \neq m} \left(v^k(S_m^k) - v^k(\mathcal{N}) \right)$. Thus, (6.29) and (6.30) imply $\hat{u}^m \leq u^m$ in case (b) as well.

We conclude that after any history, provided that her opponents bid truthfully, a bidder

maximizes her payoff by bidding truthfully. Observe that if bidders truthfully reveal their demand, at all steps (starting from H_0), we have $\kappa_0^m = \sum_{k \neq m} (-p^k(1, S^k(1)) + v^k(S^k(1)))$, and $\{S_l^m\}$ is the efficient allocation for every auction $l \in \{0, \dots, M\}$. On the other hand, since initially all prices are equal to zero, and bidders demand all items \mathcal{N} , we obtain $\kappa_0^m = \sum_{k \neq m} v^k(\mathcal{N})$, and (6.28) implies that $\Gamma^m = \gamma^m(\{S^m(T_0)\} | \{v^k\}_{k \neq m})$, where γ^m denotes the VCG payment of bidder m (see Definition 6.3.3). Therefore, the final allocation $\{S_0^m\}$ is efficient, and the corresponding payments are VCG payments, and the result follows. \square

Proof of Theorem 6.5.1. (i) Observe that if the auction terminates, then the conditions of Step S2 hold. Let $\{S_0^m\}$, and $\{S_k^m\}$ be as defined in this step. Since bidders report their demand truthfully, this implies that π^m, π^s given in the theorem statement can alternatively be expressed as:

$$\pi^m = v^m(S) - \sum_{i \in S} p_i^m - \sum_{ij \in S} p_{ij}^m, \quad (6.31)$$

for all $S \in D^m$, and

$$\pi^s = \sum_m p^m(S_0^m). \quad (6.32)$$

Observe that Step S4 implies that $p_i^m \geq 0$. Together with the construction of π^m, π^s variables, this implies that $(\{p_i^m\}, \{p_{ij}^m\}, \{\pi^m\}, \pi^s)$ is feasible in D5. Additionally, it can be checked that this solution satisfies complementary slackness conditions with a primal feasible solution of LP5 $(\{x^m\}, \{\delta_\mu\})$, such that for every m $x^m(S_0^m) = 1$, $x^m(S) = 0$ for $S \neq S_0^m$, and $\delta_{\{S_0^m\}} = 1$, $\delta_\mu = 0$ for $\mu \neq \{S_0^m\}$. Thus, it follows that $(\{p_i^m\}, \{p_{ij}^m\}, \{\pi^m\}, \pi^s)$ is optimal in D5.

Consider a formulation of D5 with bidders $\mathcal{M} - \{k\}$. Observe that the restriction of prices and bidder surpluses of $(\{p_i^m\}, \{p_{ij}^m\}, \{\pi^m\}, \pi^s)$ to $\mathcal{M} - \{k\}$ satisfies constraints of D5 involving π^m variables, since this solution satisfies same conditions in a formulation of D5 with bidders \mathcal{M} . On the other hand, Step S2 of the auction implies that $\pi_k^s \triangleq \sum_{m \neq k} p^m(S_k^m) \geq \sum_{m \neq k} p^m(\hat{S}_k^m)$ for any other complete allocation $\{\hat{S}_k^m\}_{m \neq k}$ of items to bidders. Thus, it follows that $(\{p_i^m\}_{m \in \mathcal{M} - \{k\}}, \{p_{ij}^m\}_{m \in \mathcal{M} - \{k\}}, \{\pi^m\}_{m \in \mathcal{M} - \{k\}}, \pi_k^s)$ is feasible in a formulation of D5, with bidders $m \in \mathcal{M} - \{k\}$. In addition, it can be checked that this solution satisfies complementary slackness condition with the primal feasible solution $(\{x^m\}_{m \in \mathcal{M} - \{k\}}, \{\delta_\mu\})$ such that for every $m \neq k$, $x^m(S_k^m) = 1$, $x^m(S) = 0$ for $S \neq S_k^m$, and $\delta_{\{S_k^m\}} = 1$, $\delta_\mu = 0$ for $\mu \neq \{S_k^m\}$. This implies that the aforementioned solution is also optimal in a formulation of D5 with bidders $m \in \mathcal{M} - \{k\}$.

Since a restriction of prices and bidder surpluses of $(\{p_i^m\}, \{p_{ij}^m\}, \{\pi^m\}, \pi^s)$ to bidders $m \in \mathcal{M} - \{k\}$, agrees with an optimal solution of a formulation of D5 with bidders $m \in \mathcal{M} - \{k\}$, and k is arbitrary, we conclude that this dual solution is a special optimal solution of D5.

(ii) We will establish the ex-post perfect equilibrium result by showing that after any history H_t , the second condition of Theorem 6.3.1 holds, for the truthful bidding strategy.

In obtaining a proof of this part, we will make use of some auxiliary lemmas. The proofs of these lemmas can be found after the proof of the theorem.

We first show that if after H_t bidder m bids truthfully, then the condition of Step S4 can hold at most once for $r > t$ for bidder m .

Assume that at \bar{t} this condition holds, and for bidder m the variables p_i^m, p_{ij}^m, Ψ^m are set as suggested in Step S4. We claim that after \bar{t} , we have $p_i^m \geq w_i^m$, for all i , and $p_{ij}^m \geq w_{ij}^m$ for all $(i, j) \in E$.

Lemma 6.8.1. *Assume that at time \bar{t} a condition of Step 4 holds, and p_i^m, p_{ij}^m, Ψ^m are updated accordingly. If bidder m bids truthfully after \bar{t} , then at all times $r > \bar{t}$ we have $p_i^m \geq w_i^m$, for all i , and $p_{ij}^m \geq w_{ij}^m$ for all $(i, j) \in E$.*

This lemma implies that $p_i^m \geq w_i^m \geq 0$ after \bar{t} . Additionally, it suggests that if at some time $r > \bar{t}$ we have $S \in D^m$, then it should be the case that $p_i^m = w_i^m$ and $p_{ij}^m = w_{ij}^m$. On the other hand, this implies that $S' \in D^m$ for all $S' \subset S$. These suggest that conditions of Step S4 cannot hold after \bar{t} . Thus, we conclude if after H_t bidder m bids truthfully, then the condition of Step S4 can hold at most once for $r > t$ for bidder m . Since this is true for all bidders, it follows that conditions of Step S4 do not hold for any bidder after some $\hat{t} \geq \bar{t}$.

Second, we show that if all bidders bid truthfully, and Step S4 does not hold after time \hat{t} for any bidder, then at some $r \geq \hat{t}$, the conditions of Step S2 hold.

Lemma 6.8.2. *Assume that all bidders bid truthfully after \hat{t} . Then, at some time $r \geq \hat{t}$, the conditions of Step S2 holds.*

Thus, if after any history H_t all bidders m bid truthfully, the auction terminates. On the other hand, part (i) implies that the prices that emerge when the auction terminates are special optimal prices. It can be checked that the primal feasible solution $x^m(S_0^m) = 1$, $\delta_{\{S_0^m\}} = 1$, and $x^m(S) = \delta_\mu = 0$ for remaining m, S, μ satisfies complementary slackness conditions with this dual solution. Thus, the allocation that the auction obtains at the end, $\{S_0^m\}$, is the efficient allocation.

Next assume that bidder m uses strategy z^m , after history H_t , and the auction terminates by assigning bundles S^m to bidder m . The proof of part (i) suggests that when bidders $k \neq m$ bid truthfully $p^k(S_m^k) - p^k(S_0^k) = v^k(S_m^k) - v^k(S_0^k)$. Thus, bidder m 's final payment is equal to $\sum_{k \neq m} p^k(S_m^k) - p^k(S_0^k) = \sum_{k \neq m} v^k(S_m^k) - v^k(S_0^k) = \hat{\gamma}(S^m, \{v^k\}_{k \neq m})$.

This implies that the second condition of Theorem 6.3.1 holds after any history H_t . Thus, truthful bidding is an ex-post perfect equilibrium strategy. Additionally, when bidders bid truthfully the final allocation is efficient, and hence payments $\hat{\gamma}(S^m, \{v^k\}_{k \neq m})$ correspond to VCG payments, and the claim follows. \square

Proof of Lemma 6.8.1. Observe that the claim holds at $\bar{t} + 1$, since at \bar{t} prices are updated as in Step S4. Assume that it holds until time $r \geq \bar{t}$. We will provide an inductive proof

for the claim by showing that the claim holds at $r + 1$ as well.

At time r the prices are updated either as suggested in Step S3, or as in Step S4. If it is the latter, it immediately follows that at $r + 1$ we have $p_i^m \geq w_i^m$, for all i , and $p_{ij}^m \geq w_{ij}^m$ for all $(i, j) \in E$.

Assume that at r prices are updated as suggested in Step S3. Observe that if $p_i^m > w_i^m$ then after price update $p_i^m \geq w_i^m$. On the other hand if $p_i^m = w_i^m$, then $\{i\} \in D^m$ (since $p^m(S) \geq v^m(S)$ at time r) and hence p_i^m is not updated. Thus, we conclude that for node prices $p_i^m \geq w_i^m$ after \bar{t} .

Consider an edge (i, j) such that $i, j \in \Psi^m$, and $\{i, j\} \notin D^m$. This implies that $\{i\}, \{j\} \in D^m$ before r . Since until time r the claim holds and $p^m(S) \geq v^m(S)$, it follows that $p_i^m = w_i^m$ and $p_j^m = w_j^m$. Moreover, since $\{i, j\} \notin D^m$, we have $p_{ij}^m > w_{ij}^m$. Thus, after the price update $p_{ij}^m \geq w_{ij}^m$ for edge prices as well.

Hence, it follows that $p_i^m \geq w_i^m$ for all i , and $p_{ij}^m \geq w_{ij}^m$ for all $(i, j) \in E$ holds at $r + 1$ as well. By induction, we establish the claim. \square

Proof of Lemma 6.8.2. Consider some bidder m , and observe that at $r > \hat{t}$ if there exists some i such that $\{i\} \notin D^m$, then p_i^m is decreased. Conversely assume that at time r , for all i we have $\{i\} \in D^m$. Then at time $r + 1$, we have $\Psi^m = \mathcal{N}$. Consider any edge (i, j) after $r + 1$. If $\{i, j\} \notin D^m$, then p_{ij}^m decreases. Since prices are lower bounded by weights, as suggested by Lemma 6.8.1, it follows that if conditions of Step S2 do not hold, then eventually, we have $\{i\} \in D^m$, and $\{i, j\} \in D^m$ for all $i \in \mathcal{N}$, and $(i, j) \in E$. Note that by Lemma 6.8.1 this suggest that $p_i^m = w_i^m$, and $p_{ij}^m = w_{ij}^m$ for all $i \in \mathcal{N}$, and $(i, j) \in E$. Hence, it follows that $D^m = 2^{\mathcal{N}}$.

Since this is true for all bidders, it follows that if conditions of Step S2 do not hold, eventually $D^m = 2^{\mathcal{N}}$, for all bidders. On the other hand, it can be easily checked that the conditions of Step S2 holds in this case. Hence, the claim follows. \square

Chapter 7

Conclusions

In this chapter, we conclude the thesis by providing an overview of its main contributions, and outlining some interesting future directions.

7.1 Summary

In the first part of the thesis we provided a decomposition of space of games into three orthogonal components: potential, harmonic, and nonstrategic. The first component captures the strategic properties of potential games, a well-studied class of games with desirable equilibrium and dynamic properties. The second component is identified in our work, and leads to a class of games with qualitatively very different properties from potential games. The last component, on the other hand, does not have an impact on any of the strategic properties of a given game, but it determines the efficiency of equilibria. We established that the decomposition is a useful tool for approximating a given game with a potential game, and characterizing its equilibria through this approximation. We also showed that the approximation can be valuable for analyzing the convergence of adaptive dynamics in multi-agent settings. In particular, we provided a characterization of outcome of better/best response dynamics, logit-response dynamics, and fictitious play in an arbitrary game, by exploiting properties of nearby potential games. The results of this part of the thesis both reveal some fundamental topological structures of games, and show how these can be exploited for characterization of equilibria and outcome of dynamics in various game-theoretic settings.

In the second part of the thesis, we focused on developing iterative auction formats that guarantee efficiency in multi-item settings. Since this problem is intractable in general, we restricted our attention to a special class of valuation functions, where preferences of agents can be expressed by using a graphical model. We obtained different linear programming formulations for finding the efficient outcome, and exploited algorithms that iteratively solve these linear problems for the design of iterative auctions. In particular, following this approach, we provided iterative auction formats that rely on an anonymous item pricing rule

to implement the efficient outcome when the underlying value graph has a tree structure (and satisfies a sign consistency condition). Additionally, we showed that for general graphical valuations it may be necessary to employ a more general pricing rule (bidder-specific graphical pricing) in order to implement the efficient outcome. Accordingly, we designed new iterative auction formats that rely on bidder-specific graphical pricing, and guarantee efficiency for all graphical valuation functions. Our results in this part of the thesis provide means of developing simple iterative auction formats by exploiting the special structure of graphical valuations.

7.2 Future Research: Decomposition & Dynamics

We next outline some interesting research directions related to the first part of the thesis.

Dynamics in harmonic games: While the behavior of multi-agent dynamics in potential games is reasonably well-understood, there seems to be a number of interesting research questions regarding their harmonic counterpart. In (Candogan et al., 2011c), we made some partial progress in this direction, by showing that in harmonic games, the uniformly mixed strategy profile is the unique equilibrium of the continuous time fictitious-play dynamics, and that this equilibrium point is locally stable. The global stability of this equilibrium in two-player harmonic games (which are also zero-sum games) follows from known results on the convergence of fictitious play dynamics in zero-sum games. Analyzing global stability of equilibrium under fictitious play dynamics for harmonic games with $M > 2$ players and convergence of other adaptive dynamics are interesting research directions.

Heterogeneous update rules: In the first part of the thesis we only analyzed dynamics in settings where all players update their strategies using the same rules. For instance, we assumed that all players adopt best response, or logit response dynamics with the same parameter. The limiting behavior of dynamic processes, where players adhere to different update rules is still an open question, even for potential games. An interesting future research question is whether the techniques in the first part of the thesis can be used to understand the limiting behavior of such update rules. For example, consider a potential game where all players update their strategies using a logit response with different but “close” τ parameters. Can the outcome of this dynamic process be approximated with the outcome of logit response in a related potential game where all players use the same parameter for their updates?

Guaranteeing desirable limiting behavior: Another promising research direction is to use our understanding of simple update rules, such as better/best response and logit response dynamics, to design mechanisms that guarantee desirable limiting behavior, such as low efficiency loss and “fair” outcomes in various game-theoretic settings. It is well known

that equilibria in games can be very different in terms of such properties (Roughgarden, 2005). Hence, it is of interest to develop mechanisms that ensure that strategy updates of agents converge to an equilibrium with desirable properties. It has been shown that in some cases simple pricing mechanisms can ensure convergence to desirable equilibria in near-potential games (Candogan et al., 2010b). It is an interesting research direction to extend such mechanisms to general game-theoretic settings.

7.3 Future Research: Graphical Valuations and Mechanism Design

An outline of future research directions, related to the second part of the thesis is provided below.

Robustness of iterative auctions: The results provided in the second part of this thesis rely on the assumption that the valuations of bidders can be modeled by graphical valuations. However, in practical settings, we expect graphical valuations to be only approximations of reality. How sensitive are the results presented in this thesis to the deviations from the graphical valuation assumption? For instance, do the auction formats we provide lead to inefficiency, if the true valuations of the bidders are not graphical valuations, but are approximated by graphical valuations? If so, is it possible to provide bounds on the resulting inefficiency? What are the qualitative properties of valuation functions for which the auctions we provide achieve approximate efficiency?

Auctions with value externalities: It is interesting to see if the results of this thesis can be extended to settings that allow for value externalities between bidders. For instance, assume that the value a bidder has for the set of items she acquires, not only depends on the items she receives, but also on the items that are acquired by her opponents.¹ In this case, can we formulate the efficient allocation problem as a linear program, and develop iterative auctions following the approach presented in this thesis?

Interdependent valuations: It is known that in single-item settings where valuations of bidders are interdependent, iterative auctions have interesting revenue properties. In particular, for such valuation functions, single-item iterative auctions may lead to higher revenues for the seller than the sealed bid alternatives, while preserving efficiency (Krishna, 2009). Do similar conclusions hold for multi-item auctions? A simple class of valuation functions for which this question can be studied, is the class of graphical valuation. It is an interesting future direction to understand revenue properties of various auction formats, in settings where valuations of bidders are graphical.

¹Such value externalities are commonly found in settings where bidders are competitors in different markets. A notable example is the case of patent auctions.

Empirical work: For what type of auction environments are graphical valuations a good approximation of the true valuations of bidders? How do the auctions, introduced in this thesis, compare with other multi-item auctions employed in such environments? These questions require a rigorous empirical analysis. We believe that this is an interesting direction for future work.

Special preference structures and market design: Graphical models, and other models that capture special preference structures of agents, provide us with a compact representation of complex systems. As we demonstrate in this thesis, these compact representations can be employed together with tools from operations research and game theory to systematically design simple and improved mechanisms. What are the scope and limitations of this approach? Can we rely on structured models to improve market design in settings other than auctions? For instance, is it possible to use similar simple preference models, in the context of matching markets, or markets with many sellers and buyers?

Bibliography

- Abraham, I., Babaioff, M., Dughmi, S., Roughgarden, T., 2012. Combinatorial auctions with restricted complements. In: Proceedings of the 13th ACM Conference on Electronic Commerce. ACM, pp. 3–16.
- Ahn, D., 2006. A generic bound on cycles in two-player games. Mimeo.
- Alós-Ferrer, C., Netzer, N., 2010. The logit-response dynamics. *Games and Economic Behavior* 68 (2), 413–427.
- Anantharam, V., Tsoucas, P., 1989. A proof of the Markov chain tree theorem. *Statistics & Probability Letters* 8 (2), 189–192.
- Arslan, G., Marden, J., Shamma, J., 2007. Autonomous vehicle-target assignment: A game-theoretical formulation. *Journal of Dynamic Systems, Measurement, and Control* 129 (5), 584–596.
URL <http://link.aip.org/link/?JDS/129/584/1>
- Ausubel, L., 2004. An efficient ascending-bid auction for multiple objects. *The American Economic Review* 94 (5), 1452–1475.
- Ausubel, L., 2006. An efficient dynamic auction for heterogeneous commodities. *The American Economic Review* 96 (3), 602–629.
- Ausubel, L., Cramton, P., 2004. Auctioning many divisible goods. *Journal of the European Economic Association* 2 (2-3), 480–493.
- Ausubel, L., Milgrom, P., 2002. Ascending auctions with package bidding. *The BE Journal of Theoretical Economics* 1 (1), 1–42.
- Ausubel, L. M., Milgrom, P., 2006. The lovely but lonely Vickrey auction. In: Cramton, P., Shoham, Y., Steinberg, R. (Eds.), *Combinatorial Auctions*. MIT Press, pp. 17–40.
- Başar, T., Ho, Y., 1974. Informational properties of the Nash solutions of two stochastic nonzero-sum games. *Journal of Economic Theory* 7 (4), 370–387.

- Benaïm, M., Hofbauer, J., Sorin, S., 2005. Stochastic approximations and differential inclusions. *SIAM Journal on Control and Optimization* 44 (1), 328–348.
- Berge, C., 1963. *Topological spaces*. Oliver & Boyd, London.
- Bertsekas, D., Tsitsiklis, J., 1989. *Parallel and distributed computation*. Old Tappan, NJ (USA); Prentice Hall Inc.
- Bertsimas, D., Tsitsiklis, J., 1997. *Introduction to Linear Optimization*. Athena Scientific, Belmont, MA.
- Bikhchandani, S., de Vries, S., Schummer, J., Vohra, R., 2002. Linear programming and Vickrey auctions. In: Dietrich, B., Vohra, R. (Eds.), *Mathematics of the Internet: E-auction and Markets*. New York: Springer Verlag, pp. 75–116.
- Bikhchandani, S., de Vries, S., Schummer, J., Vohra, R., 2011. An ascending Vickrey auction for selling bases of a matroid. *Operations Research* 59 (2), 400–413.
- Bikhchandani, S., Mamer, J., 1997. Competitive equilibrium in an exchange economy with indivisibilities. *Journal of Economic Theory* 74 (2), 385–413.
- Bikhchandani, S., Ostroy, J., 2002. The package assignment model. *Journal of Economic Theory* 107 (2), 377–406.
- Blume, L., 1993. The statistical mechanics of strategic interaction. *Games and Economic Behavior* 5 (3), 387–424.
URL <http://www.sciencedirect.com/science/article/pii/S0899825683710237>
- Blume, L., 1997. Population games. In: Arthur, W., Durlauf, S., Lane, D. (Eds.), *The economy as an evolving complex system II*. Addison-Wesley, pp. 425–460.
- Blumrosen, L., Nisan, N., 2010. On the computational power of demand queries. *SIAM Journal on Computing* 39 (4), 1372.
- Brown, G., 1951. Iterative solution of games by fictitious play. *Activity analysis of production and allocation* 13 (1), 374–376.
- Candogan, O., Menache, I., Ozdaglar, A., Parrilo, P. A., 2011a. Flows and decompositions of games: Harmonic and potential games. Forthcoming in *Mathematics of Operations Research*.
URL <http://mor.journal.informs.org/cgi/content/abstract/moor.1110.0500v1>
- Candogan, O., Ozdaglar, A., Parrilo, P., December 2010a. A projection framework for near-potential games. In: *49th IEEE Conference on Decision and Control (CDC)*. pp. 244–249.

- Candogan, O., Ozdaglar, A., Parrilo, P., 2011b. Dynamics in games and near-potential games. LIDS, MIT, Tech. Rep.
- Candogan, O., Ozdaglar, A., Parrilo, P. A., 2011c. Learning in near-potential games. In: Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on. IEEE, pp. 2428–2433.
- Candogan, U. O., Menache, I., Ozdaglar, A., Parrilo, P. A., 2010b. Near-optimal power control in wireless networks: a potential game approach. In: INFOCOM, 2010 Proceedings IEEE. IEEE, pp. 1–9.
- Cho, G., Meyer, C., 2001. Comparison of perturbation bounds for the stationary distribution of a Markov chain. *Linear Algebra and its Applications* 335 (1-3), 137–150.
- Chopra, S., 1994. The graph partitioning polytope on series-parallel and 4-wheel free graphs. *SIAM journal on discrete mathematics* 7 (1), 16–31.
- Christodoulou, G., Mirrokni, V., Sidiropoulos, A., 2006. Convergence and approximation in potential games. *Lecture Notes in Computer Science* 3884, 349.
- Chung, F. R., 1997. *Spectral Graph Theory* (CBMS Regional Conference Series in Mathematics, No. 92). American Mathematical Society.
- Cramton, P., Ausubel, L., McAfee, R., McMillan, J., 1997. Synergies in wireless telephony: Evidence from the broadband pcs auctions. *Journal of Economics and Management Strategy* 6 (3), 497–527.
- Cramton, P., Shoham, Y., Steinberg, R., 2006. *Combinatorial auctions*. MIT Press.
- Cressman, R., Morrison, W., 1998. On the evolutionary dynamics of crime. *The Canadian Journal of Economics/Revue canadienne d'Economie* 31 (5), 1101–1117.
- Daskalakis, C., Goldberg, P. W., Papadimitriou, C. H., 2006. The complexity of computing a Nash equilibrium. In: *Proceedings of the thirty-eighth annual ACM symposium on theory of computing*. ACM, pp. 71–78.
- Daskalakis, K., Papadimitriou, C., 2005. The complexity of games on highly regular graphs. In: Brodal, G., Leonardi, S. (Eds.), *Algorithms ESA 2005*. Vol. 3669 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, pp. 71–82.
URL http://dx.doi.org/10.1007/11561071_9
- De Vries, S., Schummer, J., Vohra, R., 2007. On ascending Vickrey auctions for heterogeneous objects. *Journal of Economic Theory* 132 (1), 95–118.
- De Vries, S., Vohra, R., 2003. Combinatorial auctions: A survey. *INFORMS Journal on Computing* 15 (3), 284–309.

- Diestel, R., 2005. Graph Theory Graduate Texts in Mathematics; 173. Springer-Verlag Berlin and Heidelberg GmbH & Company KG.
- Dubey, P., Haimanko, O., Zapechelnyuk, A., 2006. Strategic complements and substitutes, and potential games. *Games and Economic Behavior* 54 (1), 77 – 94.
URL <http://www.sciencedirect.com/science/article/pii/S089982560400171X>
- Duffin, R., 1965. Topology of series-parallel networks. *Journal of Mathematical Analysis and Applications* 10 (2), 303–318.
- Engelbrecht-Wiggans, R., Kahn, C., 1991. Protecting the winner: Second-price versus oral auctions. *Economics Letters* 35 (3), 243–248.
- Facchini, G., van Meegen, F., Borm, P., Tijs, S., 1997. Congestion models and weighted bayesian potential games. *Theory and Decision* 42 (2), 193–206.
- Freidlin, M., Wentzell, A., 1998. Random perturbations of dynamical systems. Springer Verlag.
- Friedman, D., 1991. Evolutionary games in economics. *Econometrica* 59 (3), 637.
- Fudenberg, D., Levine, D., 1998. The theory of learning in games. MIT press.
- Fudenberg, D., Tirole, J., 1991. Game Theory. MIT Press.
- Germano, F., 2006. On some geometry and equivalence classes of normal form games. *International Journal of Game Theory* 34 (4), 561–581.
- Gilboa, I., Schmeidler, D., 1995. Canonical representation of set functions. *Mathematics of Operations Research* 20 (1), 197–212.
- Goemans, M., Mirrokni, V., Vetta, A., 2005. Sink equilibria and convergence. In: FOCS '05: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science. IEEE Computer Society, Washington, DC, USA, pp. 142–154.
- Golub, G., Van Loan, C., 1996. Matrix computations. Johns Hopkins University Press.
- Gul, F., Stacchetti, E., 1999. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory* 87 (1), 95–124.
- Gul, F., Stacchetti, E., 2000. The English auction with differentiated commodities. *Journal of Economic theory* 92 (1), 66–95.
- Hammond, P., 2005. Utility invariance in non-cooperative games. In: Advances in public economics: utility, choice and welfare. Springer, pp. 31–50.

- Haviv, M., Van der Heyden, L., 1984. Perturbation bounds for the stationary probabilities of a finite Markov chain. *Advances in Applied Probability* 16 (4), 804–818.
- Hillas, J., Kohlberg, E., 2002. Foundations of strategic equilibrium. *Handbook of game theory* 3, 1597–1663.
- Hofbauer, J., Hopkins, E., 2005. Learning in perturbed asymmetric games. *Games and Economic Behavior* 52 (1), 133–152.
- Hofbauer, J., Sandholm, W., 2002. On the global convergence of stochastic fictitious play. *Econometrica*, 2265–2294.
- Hofbauer, J., Schlag, K., 2000. Sophisticated imitation in cyclic games. *Journal of Evolutionary Economics* 10 (5), 523–543.
- Hohner, G., Rich, J., Ng, E., Reid, G., Davenport, A. J., Kalagnanam, J. R., Lee, H. S., An, C., January/February 2003. Combinatorial and quantity-discount procurement auctions benefit mars, incorporated and its suppliers. *Interfaces* 33 (1), 23–35.
URL <http://interfaces.journal.informs.org/content/33/1/23.abstract>
- Jensen, T. R., Toft, B., 1995. *Graph coloring problems*. New York, NY: John Wiley & Sons.
- Jiang, X., Lim, L.-H., Yao, Y., Ye, Y., 2011. Statistical ranking and combinatorial Hodge theory. *Mathematical Programming* 127 (1), 203–244.
- Jordan, J., 1993. Three problems in learning mixed-strategy Nash equilibria. *Games and Economic Behavior* 5 (3), 368–386.
- Kalai, A., Kalai, E., 2010. *Engineering Cooperation in Two-Player Strategic Games*. Mimeo.
- Kannan, R., Theobald, T., 2010. Games of fixed rank: a hierarchy of bimatrix games. *Economic Theory* 42 (1), 157–173.
URL <http://ideas.repec.org/a/spr/joecth/v42y2010i1p157-173.html>
- Kelso, A., Crawford, V., 1982. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society*, 1483–1504.
- Khalil, H., 2002. *Nonlinear systems*. Prentice Hall, NJ: Prentice Hall.
- Kleinberg, N., Weiss, J., 1985. Algebraic structure of games. *Mathematical Social Sciences* 9 (1), 35–44.
- Kleinberg, N., Weiss, J., 1986. The orthogonal decomposition of games and an averaging formula for the Shapley value. *Mathematics of Operations Research*, 117–124.
- Krishna, V., 2009. *Auction theory*. Academic press.

- Lehmann, B., Lehmann, D., Nisan, N., 2006. Combinatorial auctions with decreasing marginal utilities. *Games and Economic Behavior* 55 (2), 270–296.
- Leighton, F., Rivest, R., 1983. The Markov chain tree theorem. Tech. Rep. MIT/LCS/TM-249, MIT, Laboratory for Computer Science.
- Li, D., Zhang, X., 2002. On dynamical properties of general dynamical systems and differential inclusions. *Journal of Mathematical Analysis and Applications* 274 (2), 705 – 724. URL <http://www.sciencedirect.com/science/article/pii/S0022247X02003529>
- Marden, J., Arslan, G., Shamma, J., 2009a. Cooperative Control and Potential Games. *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics* 39 (6), 1393–1407.
- Marden, J., Arslan, G., Shamma, J., 2009b. Joint strategy fictitious play with inertia for potential games. *IEEE Transactions on Automatic Control* 54 (2), 208–220.
- Marden, J., Shamma, J., 2012. Revisiting log-linear learning: Asynchrony, completeness and payoff-based implementation. *Games and Economic Behavior*.
- Marinacci, M., 1996. Decomposition and representation of coalitional games. *Mathematics of Operations Research*, 1000–1015.
- McAfee, R. P., McMillan, J., Wilkie, S., 2010. The greatest auction in history. In: *Better living through economics*. Harvard University Press, pp. 168–184.
- Mertens, J., 2004. Ordinality in non cooperative games. *International Journal of Game Theory* 32 (3), 387–430.
- Mishra, D., Parkes, D., 2007. Ascending price Vickrey auctions for general valuations. *Journal of Economic Theory* 132 (1), 335–366.
- Mishra, D., Parkes, D., 2009. Multi-item Vickrey–Dutch auctions. *Games and Economic Behavior* 66 (1), 326–347.
- Monderer, D., Shapley, L., 1996a. Fictitious play property for games with identical interests. *Journal of Economic Theory* 68 (1), 258–265.
- Monderer, D., Shapley, L., 1996b. Potential games. *Games and economic behavior* 14 (1), 124–143.
- Moreton, P., Spiller, P., 1998. What’s in the air: Interlicense synergies in the federal communications commission’s broadband personal communication service spectrum auctions. *Journal of Law and Economics* 41 (2), 677–716.
- Morris, S., Ui, T., 2004. Best response equivalence. *Games and Economic Behavior* 49 (2), 260–287.

- Moulin, H., Vial, J., 1978. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory* 7 (3), 201–221.
- Neyman, A., 1997. Correlated equilibrium and potential games. *International Journal of Game Theory* 26 (2), 223–227.
- Nisan, N., Roughgarden, T., Tardos, E., Vazirani, V. V., 2007. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA.
- Nisan, N., Schapira, M., Valiant, G., Zohar, A., 2011a. Best-response auctions. In: *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, pp. 351–360.
- Nisan, N., Schapira, M., Valiant, G., Zohar, A., 2011b. Best-response mechanisms. In: *ICS11: Proceedings of the 2nd symposium on Innovations in Computer Science*.
- Nisan, N., Segal, I., 2006. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory* 129 (1), 192–224.
- Papadimitriou, C., Steiglitz, K., 1998. *Combinatorial optimization: algorithms and complexity*. Dover Publications.
- Parkes, D., 1999. i-bundle: An efficient ascending price bundle auction. In: *Proceedings of the 1st ACM conference on Electronic commerce*. ACM, pp. 148–157.
- Parkes, D., 2006. Iterative combinatorial auctions. In: Cramton, P., Shoham, Y., Steinberg, R. (Eds.), *Combinatorial Auctions*. MIT Press, pp. 41–77.
- Polthier, K., Preuß, E., 2002. Identifying vector field singularities using a discrete Hodge decomposition. *Visualization and Mathematics III*, 113–134.
- Rosenthal, R., 1974. Correlated equilibria in some classes of two-person games. *International Journal of Game Theory* 3 (3), 119–128.
- Rothkopf, M., Teisberg, T., Kahn, E., 1990. Why are Vickrey auctions rare? *Journal of Political Economy*, 94–109.
- Roughgarden, T., 2005. *Selfish Routing and the Price of Anarchy*. MIT Press.
- Sandholm, W., 2010a. *Population games and evolutionary dynamics*. MIT Press, Cambridge, MA.
- Sandholm, W. H., 2010b. Decompositions and potentials for normal form games. *Games and Economic Behavior* 70 (2), 446–456.
- Shamma, J., Arslan, G., 2004. Unified convergence proofs of continuous-time fictitious play. *IEEE Transactions on Automatic Control* 49 (7), 1137–1141.

- Shapley, L., 1953. A value for n -person games. In *Contributions to the Theory of Games*, II. ed. H. W. Kuhn and A. W. Tucker, 307–317.
- Shapley, L., 1964. Some topics in two-person games. *Advances in game theory* 52, 1–29.
- Smith, J., Hofbauer, J., 1987. The “battle of the sexes”: a genetic model with limit cycle behavior. *Theoretical population biology* 32 (1), 1–14.
- Sun, N., Yang, Z., 2006. Equilibria and indivisibilities: gross substitutes and complements. *Econometrica* 74 (5), 1385–1402.
- Sun, N., Yang, Z., 2009. A double-track adjustment process for discrete markets with substitutes and complements. *Econometrica* 77 (3), 933–952.
- Tong, Y., Lombeyda, S., Hirani, A., Desbrun, M., 2003. Discrete multiscale vector field decomposition. *ACM Transactions on Graphics* 22 (3), 445–452.
- Uno, H., 2007. Nested potential games. *Economics Bulletin* 3 (17), 1–8.
- Vohra, R. V., 2011. *Mechanism Design: A linear programming approach*. Cambridge University Press.
- Von Neumann, J., Morgenstern, O., 1947. *Theory of Games and Economic Behavior*. Princeton University Press.
- Voorneveld, M., 2000. Best-response potential games. *Economics Letters* 66 (3), 289–295.
- Voorneveld, M., Borm, P., van Meegen, F., Tijs, S., Facchini, G., 1999. Congestion games and potentials reconsidered. *International Game Theory Review* 1 (3&4), 283–299.
- Young, H., 1993. The evolution of conventions. *Econometrica: Journal of the Econometric Society* 61 (1), 57–84.
- Young, H., 2004. *Strategic learning and its limits*. Oxford University Press, USA.
- Zhou, H., Berry, R., Honig, M., Vohra, R., 2009. Complementarities in spectrum markets. In: *Communication, Control, and Computing, 2009. Allerton 2009. 47th Annual Allerton Conference on*. IEEE, pp. 499–506.