Algorithmic Approaches to Oligopoly Theory

Nithum Thain

Doctor of Philosophy

Department of Mathematics and Statistics

McGill University

Montreal, Quebec

May, 2015

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Doctor of Philosophy.

Copyright ©Nithum Thain, 2015

DEDICATION

Dedicated to my grandparents

ACKNOWLEDGEMENTS

This PhD would not exist but for the support of a great community of friends, family, and mentors.

First and foremost, I'd like to thank my supervisor Adrian Vetta, who introduced me to the field of Algorithmic Game Theory and guided me through both a Master's and a PhD. His encyclopaedic knowledge, deep analysis, and patient advice were the defining factors in my graduate school success. I am lucky to have had him as a supervisor.

I would also like to thank my family, Dev, Mom, and Dad. Without you three to keep me sane, I'm not sure how I would have managed the hecticness of these past few years. Thanks also to my girlfriend, Sulin, who patiently put up with the highs, lows, and uncertainties of this whole process. Her love and support have value beyond measure.

Much of my research was performed in collaboration, and I would like to thank my collaborators Ruta Mehta, László Végh, and Vahab Mirrokni. Their insights led to crucial steps forward in the research program. I would also like to thank my academic mentors, Bruce Shepherd, Vašek Chvátal, Mike Roth, Georg Gottlob, and Ram Murty for introducing me to the beauty of so many fields of mathematics.

A PhD is a lonely endeavour and having a few good friends around made a difference. Here I'll name just a few of the many people whose friendship helped ease the way. My sincerest thanks to: Justin So, Mike White, Erich Chung, Derrick Lee, Tzu

Hao Hsu, Ryan Nelson, Jonathan Gerber, Emily Mawe, Steven Maye, Sue Humphrey, Rosanna Nicol, Neil Olver, Tyler Grant, Julien Gagnon, Gillian Langor, Bhaskar Bhushan, Kamal Wood, Kayli Johnson, Jaspreet Khangura, Nicholas Chesterley, Matt Lewis, Madu Jayatunga, Richard Lau, Richard Stebbing, Gaelle Coullon, Dups Wijayawardhana, CJ Adams, Léticia Villeneuve, Gabe Hughes, Kelly O'Connor, Sheiry Dhillon, Ramya Ravishankar, Aunnie Patton, Evelyn Chan, Sean Mullins, Melissa Bailey, Aneil Jaswal, Josh Chauvin, Richard Higgins, Brad Allen, Zoe Whitton, James Tilbuty, Anaar Kara, Brett House, Ryan D'Eon, Shama Acharya, Dougal Graham, Mike Leblanc, Kelsey Murrell, David Damberger, Jimmy Bray, Matthew Grenning, Tanya Taylor, Henry Barmeier, Andreas Heger, Pieter Koornhof, Brian Coulter, and Andrew Critch.

Finally, I would like to thank NSERC, the Rhodes Trust, the Schulich family for their generous support of my research through various grants and scholarships. Thanks also to Julien Gagnon for helping to translate my abstract into French.

CONTRIBUTION OF AUTHORS

The content of Chapter 3 is based on joint work with Adrian Vetta. The content of Chapter 4 is based on joint work with Vahab Mirrokni and Adrian Vetta. The content of Chapter 5 is based on joint work with Ruta Mehta, László Végh, and Adrian Vetta. The sections of these work included are those for which I was the primary researcher. The work in the remaining chapters is my own.

Copyright Notice

The work in Chapter 3 is based on an earlier work [167]: Computational aspects of multimarket price wars, in *Proceedings of 5th International Workshop on Internet and Network Economics* (WINE), ©Springer-Verlag Berlin Heidelberg 2009.

The work in Chapter 4 is based on an earlier work [116]: A theoretical examination of practical game playing: lookahead search, in *Proceedings of the 5th international conference on Algorithmic Game Theory* (SAGT), ©Springer-Verlag Berlin Heidelberg 2012.

The work in Chapter 5 is based on an earlier work [113]: To save or not to save: the Fisher game, in *Proceedings of 11th International Workshop on Internet and Network Economics* (WINE), ©Springer-Verlag Berlin Heidelberg 2014.

ABSTRACT

The computational tractability of economic games has been a core theme of the study of Algorithmic Game Theory. The underlying rationale for this research is that if equilibria cannot be computed by an algorithm in a reasonable time (usually polynomial time), then it is difficult to imagine that rational agents will be able to find it through strategic interaction. In this thesis we apply this principle to the decisions faced by market oligopolists. We study the computational complexity of decisions in oligopoly models, the effectiveness of heuristic search algorithms they can apply, and the overall social impact of strategization on imperfectly competitive markets.

We first consider the complexity of decision making with regards to predatory pricing in multimarket oligopoly models. Specifically, we present multimarket extensions of the classical single-market models of Bertrand, Cournot and Stackelberg, and introduce the War Chest Minimization Problem. This is the natural problem of deciding whether a firm has a sufficiently large war chest to win a price war. On the negative side we show that, even with complete information, it is hard to obtain any multiplicative approximation guarantee for this problem. Moreover, these hardness results hold even in the simple case of linear demand, price, and cost functions. On the other hand, we give algorithms with arbitrarily small additive approximation guarantees for the Bertrand and Stackelberg multimarket models with linear demand, price, and cost functions. Furthermore, in the absence of fixed costs, this problem is solvable in polynomial time in all our models.

We next turn our attention to the Lookahead heuristic, one of the most widely used game-playing heuristics. Given the practical importance of the method, we compare the performance of lookahead search in several economic settings: those of Cournot competition and Adword auctions. The main question we try to answer with this framework is the impact on social quality of outcome when agents apply lookahead search. Myopic game playing, where each player can only foresee the immediate effect of her own actions, is a special case of lookahead search. Thus, it is natural to ask whether social outcomes improve when players use more foresight than in myopic behaviour. We demonstrate that the answer depends on the game played.

Finally, we examine the Fisher market model when buyers, as well as sellers, have an intrinsic value for money. We show that when the buyers have oligopsonistic power they are highly incentivized to act strategically with their monetary allocations, as their potential gains are unbounded. This is in contrast to the bounded gains that have been shown when agents strategically report utilities. Our main focus is upon the consequences for social welfare when the buyers act strategically. To this end, we define the *Price of Imperfect Competition (PoIC)* as the worst case ratio of the welfare at a Nash equilibrium in the induced game compared to the welfare at a Walrasian equilibrium. We prove that the PoIC is at least 1/2 in some markets with CES utilities. Furthermore, for linear utility functions, we prove that the PoIC increases as the level of competition in the market increases. Additionally, we prove that a Nash equilibrium exists in the case of Cobb-Douglas utilities. In contrast, we show that Nash equilibria need not exist for linear utilities. However, in that case, good welfare guarantees are still obtained for the Nash dynamics of the game.

ABRÉGÉ

La résolubilité computationnelle des jeux économiques est un thème central à la théorie des jeux algorithmique. La logique sous-jacente à cette étude repose sur l'idée qu'il est difficile d'imaginer que des agents rationnels puissent parvenir aux équilibres d'un jeu dans le cadre d'interactions stratégiques si ces équilibres ne peuvent être calculés par un algorithme en un temps raisonnable. Dans cette thèse, nous appliquons ce principe aux décisions auxquelles font face les oligopoles. Nous étudions la complexité computationnelle des décisions dans les modèles d'oligopoles, l'efficacité d'algorithmes de recherche euristique qu'ils peuvent employer, et l'impact social général de l'élaboration de leurs stratégies sur les marchés imparfaitement compétitifs. Nous nous penchons en premier lieu sur la complexité décisionnelle de l'adoption prédatrice de prix dans le cadre de modèles d'oligopoles multi-marchés. Plus particulièrement, nous présentons des extensions multi-marchés des modèles classiques de Cournot, Bertrand et Stackelberg, et nous introduisons le Problème de Minimisation du Trésor de Guerre. Ce problème est celui d'une firme ayant à évaluer si elle détient un "trésor de guerre" suffisant pour remporter une guerre de prix. Nous démontrons que même avec information complète, il est difficile d'obtenir une garantie d'approximation multiplicative pour ce problème. Ce résultat tient même dans le cas simple de demande, prix, et fonctions de coûts linéaires. D'autre part, nous concevons des algorithmes avec garanties d'approximation additive arbitrairement petite pour les modèle multi-marchés de Bertrand et Stackelberg avec demande, prix et fonctions de coût linéaires. En l'absence de coûts fixes, ce problème

est solvable en temps polynomial dans tous nos modèles.

Nous abordons ensuite l'euristique "Lookahead", l'un des plus généralement utilisés dans le cadre de jeux. Étant donné l'importance pratique de cette méthode, nous comparons sa performance dans différents environnements économiques, plus particulièrement dans les situations de compétition à la Cournot et dans les enchères Adword. L'objectif principal de cette démarche est l'étude de l'impact de l'emploi d'une recherche de type lookahead par les agents sur le bien-être agrégé. Jouer de facon myopique, c'est-à-dire quand les agents jouent en ne prévoyant que les effets immédiats de leurs actions, est un cas spécial de recherche de type lookahead. Dans cet ordre d'idées, il est naturel de se demander si les résultats sociaux (e.g. production et bien-être) s'améliorent quand les agents adoptent un comportement plus prévoyant que myopique. Nous montrons que la réponse à cette question dépend du jeu.

Finalement, nous examinons le modèle de marché de Fisher quand les acheteurs et vendeurs attribuent une valeur intrinsèque à l'argent. Nous montrons que lorsques les acheteurs ont un pouvoir oligopsonique, ils ont hautement intérêt à agir de facon stratégique avec leur part d'argent puisque leurs gains potentiels sont illimités. Ce résultat contraste avec celui de gains limités obtenu lorsque les agents dévoilent leur utilité stratégiquement. Nous portons notre attention principalement sur les conséquences, en termes de bien-être social, du comportement stratégique des acheteurs. À cette fin, nous définissons le Prix de compétition imparfaite come le pire ratio possible de bien-être à un équilibre de Nash dans le jeu induit au bien-être à un équilibre Walrasien. Nous prouvons que le Prix de compétition imparfaite est au d'au moins

1/2 dans les marchés avec des fonctions d'utilité de type CES. De plus, nous prouvons que le Prix de compétition imparfaite s'accroît avec le niveau de compétition. Nous prouvons également qu'un équilibre de Nash existe toujours dans le cas des fonctions d'utilité de type Cobb-Douglas, mais pas nécessairement dans le cas de fonctions d'utilité linéaires. Dans ce dernier cas, toutefois, des garanties de haut bien-être sont obtenues pour les dynamiques de Nash du jeu.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

CHAPTER 1
Introduction

1.1 Overview

This thesis studies, from an algorithmic perspective, the challenging decisions that market participants are faced with when they find they have power in a particular marketplace. Our primary focus is on the cases of oligopoly (and oligopsony) where a few sellers (resp. buyers) have market power and can use it to affect the price they achieve on the market.

The computational tractability of economic games has been a core theme of the study of algorithmic game theory. The underlying rationale for this research is that if equilibria cannot be computed by an algorithm in a reasonable time (usually polynomial time), then it is difficult to imagine that rational agents will be able to find it through strategic interaction. To quote Kamal Jain, "If your laptop cannot find it, neither can the market" [131]. In this thesis we apply this principle to the decisions faced by market oligopolists. In Chapter 3 we formulate models of oligopolistic price wars which prove to be computationally intractable, though we are able to provide some positive approximation results. In Chapter 4 we discuss how limits on the computational ability of market agents can affect the quality of the resultant equilibria. Finally, in Chapter 5 we explore how the simple act of strategy can affect the welfare of a marketplace.

For the remainder of this chapter we examine the historical contexts of oligopoly theory and strategic market games. We introduce the fundamental models of Cournot and Bertrand which together form the basis of key subsequent work on oligopoly competition. We outline some of the key historical insights in the development of these models and current directions of research. We also discuss the roots of the algorithmic study of markets and the recent breakthroughs that have been made in the computation of market equilibria.

In Chapter 2, we go over the formal definitions of economic models for oligopoly and perfect competition. We describe the Cournot model and present equilibria for a few simple Cournot markets. We contrast this to the results that are obtained under the Stackelberg and Bertrand models. For perfect competition, we introduce the Fisher model of general equilibrium theory and define the Eisenberg-Gale convex program for computing equilibrium prices and allocations.

In Chapter 3 we study the complexity of predatory pricing in an oligopolistic multimarket setting. We introduce multimarket models of oligopoly competition and the War Chest Minimization Problem which asks how much of a War Chest an oligopolist needs to successfuly price a competitor out of the market. We demonstrate that this is an NP-hard problem to solve, even in the linear case, and that it is even NP-hard to find an approximation algorithm. However, we demonstrate polynomial time algorithms for solving the problem if there are no fixed costs. We also prove additive approximation guarantees for the Bertrand and Stackelberg models.

In Chapter 4, we study the framework of lookahead search as it applies to economic games. This framework provides heuristic algorithms that market participants can use if they have limited computing capabilites. We demonstrate that under the Cournot model of equilibria, looking ahead encourages oligopolists to increase production and thereby increase the social surplus. This effect is also shown to be non-linear with the highest social surplus achieved for 2 degrees of lookahead. We also examine Generalised Second Price Auctions, a popular online auction methodology. We show that, depending on the nuances of the lookahead model in question, the social utility that arises with 2 degrees of lookahead in this setting is either optimal or within a constant factor of optimal.

Chapter 5 explores the social cost of strategic decision making and the connection between oligopsony theory and general equilibrium theory. We introduce a strategic variant of the Fisher market called the Fisher Game. This game, based on the earlier work of Codognato, Gabszewicz, and Michel ([41],[70]), adapts Fisher's perfectly competitive context by allowing market participants some indirect influence on price. We introduce a welfare concept called the Price of Imperfect Competition which captures the social utility lost by this pricing power and demonstrate that this loss is bounded by a factor of 2. We also explain how this welfare ratio changes as market competition increases and examine the dynamic context when equilibria don't exist.

We conclude in Chapter 6 by summarizing our contributions in this thesis and articulating research directions that stem from this work.

1.2 Historical Context

In this section we review the relevant historical context of oligopoly theory and the application of algorithmic techniques to markets. We leave detailed discussions of the modern literature related to each particular topic that we study to the introductory sections of Chapters 3, 4, and 5. Our aim here is instead to describe the historical and empirical importance of both the economics of imperfect competition and the application of algorithmic techniques to market problems.

1.2.1 The History of Oligopoly

Oligopoly theory is the study of markets and industries dominated by a few sellers. Oligopolistic markets sit between the extremes of perfect competition, also known as general equilibrium theory where any individual market participant has little influence on the price of goods, and monopoly competition, where one participant can entirely dictate the price. In the oligopolistic setting, the centralization of market power allows the market participants to have a strong but incomplete influence on market prices for the goods they sell. This adds a layer of complexity not seen in either of the extreme cases, as participants must now consider not only their own production and sale decisions, but must also incorporate the decisions of other market participants. Put more simply, oligopolistic competition, unlike monopoly or perfect competition, is a game.

Antoine Augustin Cournot is credited as the founder of Oligopoly theory [119]. His seminal book *Researches on the Mathematical Principles of the Theory of Wealth*, published in 1838, introduces the first mathematical models of duopoly[1] competition in the context of the mineral water industry. Indeed the book is credited as being one of the earliest works to articulate a mathematical model in economics [135]. Cournot introduces the concept of equilibrium analysis with his *Cournot equilibrium* which anticipates the more general approach of Nash equilibrium analysis [119], a concept that John Nash would not formalize until 1949. This book also introduces models of monopoly that are still in use today and discusses key questions in the theory of perfect competition that inspired "the Marginal Revolution" in economics, an attempt to understand economic behaviour from the point of view of marginal utility.

Cournot's model examines two oligopolistic sellers competing in a market who each strategize over the quantity of good they will produce. The goods are assumed to be homogeneous, meaning that there is no way for buyers to differentiate between goods. The quantity of goods produced and the market demand function thus completely determine the price at which goods are sold. Cournot's analysis, which we will examine in more detail in Chapter 2, illustrates how to compute an equilibrium level of production where no seller has an incentive to change their production. This equilibrium has many properties considered to be realistic based on observed evidence. In particular, it allows for firms with different costs of production to coexist

[1] A duopoly is an oligopoly with only two sellers.

in the marketplace and it allows oligopolistic sellers to use their market power to generate economic profit by pricing above their marginal cost.

One of Cournot's harshest critics was Joseph Louis Francois Bertrand [18]. Bertrand challenged the concept of mathematical economics in general and was critical, in particular, of the quantity competition that underlay Cournot's work. He argued that, in the market, firms choose prices and consumers generate demand depending on this price, in stark contrast to Cournot's model. This led to the development of the Bertrand model of price competition. In this model, two oligopolists selling homogeneous goods now set a price for these goods. The demand of the market depends on this price but the market then chooses to buy entirely from the seller with the lowest price. This quickly leads to only one equilibrium outcome, whereby the seller with the lowest marginal cost of production sells at the cost of production of his most efficient rival. The implications of this model are that only sellers with the lowest marginal cost can exist in a market and that if even two of these sellers coexist, neither can make economic profit. This leads to questions about why either participant would enter the market in the presence of even the mildest fixed costs. This peculiar implication of the equilibrium is known as *Bertrand's Paradox*.

These two models of quantity and price competition laid the groundwork for over a century of subsequent oligopoly theory. It became standard to refer to all quantity based models of competition as Cournot competition and price based models as Bertrand competition. It can be argued that all subsequent models of oligopoly theory simply introduce layers of realism and complexity to these models in order to better explain observable phenomena.

An important adaptation of Cournot's model was introduced by Heinrich von Stackelberg in 1934 [162]. Stackelberg did not believe in the Cournot concept of equilibrium as it only held in the case of simultaneous action by both players. He believed that a sequential game made more sense in real world contexts. Thus, he introduced a framework for leader-follower behaviour where some firm (or set of firms) called leaders initially produce a quantity of goods and other firms, called followers, must adapt their quantity based on the leaders' production. This model allows leaders to capture more profit than in the Cournot case, at the expense of followers. This theory led American and European economists in the 1980s to recommend that the government pay export subsidies to domestic firms in certain industries to allow them to become market leaders [127].

In the realm of price competition, Edgeworth gives an early enhancement of Bertrand's model of price competition by introducing capacity constraints for producers in his 1897 paper *Me teoria pura del monopolio* [51]. His model remediates some of the challenges of Bertrand's paradox as it allows for sustainable prices above the level of perfect competition. However, in introducing capacity constraints, Edgeworth introduces discontinuities into the model that mean equilibria are no longer guaranteed to exist.

Another generalization of Bertrand's model was introduced by Edward Chamberlin [35]. Chamberlin's model introduces differentiation of goods, allowing consumers to have a preference for goods which act as substitutes and allowing for the demand of each good to depend on the price of all goods in the markets. This model and subsequent work on price based competition of differentiated goods is now studied as the Theory of Monopolistic Competition and Chamberlin is credited as its founder.

Differentiation of another type is explored by the work of Hotelling [84]. In this model, the differentiation of goods is based on location. In the original paper, Hotelling presents an example of two price-setting oligopolists selling an otherwise homogenous good who are given the option to differentiate based on location. Buyers are distributed uniformly along a line segment and the sellers must choose both the price of sale and where to locate themselves. The cost of the good to buyers is a function of both the price set by the oligopolists and the distance to the seller (transportation cost). The surprising outcome of Hotelling's model is that even given the choice to differentiate on location, the only equilibrium strategy is for both firms to be collocated at the middle of the segment. This additionally causes quite a loss in social welfare as compared to a scenario where the firms locate at the quartiles of the segment and thereby reduce transportation costs for buyers.

Lerner and Singer helped to resolve this somewhat paradoxical outcome by pointing out that Hotelling's model relied on unrealistic assumptions of inelastic demand [109]. In the original model, regardless of prices, buyers would always demand a unit of the good. Lerner and Singer and later Smithies presented models with elastic demand functions and demonstrated more reasonable equilibria for the game.

The original Bertrand and Cournot models were formulated as simultaneous, one-shot games. By studying oligopoly competition in a repeated or dyanmic setting (see for example [168] or [145]), the nature of equilibria change. It can be argued that in a repeated oligopoly game, players are incentivized to collude and form cartels so as to avoid a price war. This is yet another way that Bertrand's paradox can be resolved. In Chapter 3, we will discuss the dynamics of cartel's and price wars in more detail.

Modern oligopoly research continues to build on the complexity and realism of these models. The study of theoretical oligopoly continues in myriad directions from work on oligopoly dynamics [135] to strategies of differentiation [86] or work on capacity precommitment [107]. The theory has also found applications to fields including futures markets [4], international trade [27], electricity distribution [178], and the airline industry [28].

1.2.2 Algorithms and Markets

The study of markets from an algorithmic perspective begins in a different part of industrial organization, namely with general equilibrium theory: the study of perfect competition.

The foundation of general equilibrium theory is attributed to Léon Walras, one of the key figures in the "Marginal Revolution" of Economics. He sought to study the problem formulated by Cournot of whether supply and demand could be made to equate to clear all markets simultaneously (it was known at the time that they could easily be made to equate in a single market). In his 1874 publication *Elements of Pure Economics* [175], Walras formulated the first mathematical model of pure competition, the so-called Walras model which could model the "state of the economic system at any point in time" [7]. This model was largely ignored during Walras' lifetime but soon thereafter came to prominence, becoming one of the core models for microeconomic theory in the 20th century.[2]

The simplest version of Walras' model of general equilibrium theory starts with a market whose participants are endowed with an initial set of goods. The problem is then for the market mechanism to find prices for these goods so that players can sell and buy goods to maximize utility and so that the market clears, meaning that there is no excess supply or demand. While Walras did not establish that such *equilibrium prices* always existed, he did propose a process called tâtonnement, by which this system could converge towards an equilibrium. Walras describes the process as an auctioneer adjusting market prices based on feedback from market agents as to their demand at a given price. This process could be thought of as the first attempt at an algorithm for computing market prices.

It was not until 1954 [7] that Kenneth Arrow and Gerard Debreu were able to establish that equilibria exist in the context of pure competition. Using the Kakutani fixed point theorem, the same approach as used by John Nash in proving the existence of a mixed strategy Nash equilibrium, they demonstrated that if the participants

[2] Schumpeter, a prominent Austrian economist, acclaimed Walras as the "greatest of all economists" [176].

had concave utility functions (along with a few other technical assumptions) then an equilibrium price could be found for which the market clears. However, their proof technique was non-constructive and so did not articulate how such market prices could be found in practice.

The success of Arrow and Debreu in establishing the existence of equilibrium prices revitalized the study of how these prices could actually be computed. In 1959, Arrow, Block, and Hurwicz [6] established that Walras' tâtonnement would indeed always converge to the equilibrium prices if the goods in the market satisfied the property of weak gross substitutability (we will define this property in more detail in Section 2.2.2). Unfortunately, soon thereafter Scarf gave a simple example where tâtonnement does not converge if goods are not substitutes [143].

Scarf proposed a different approach to computing equilibria prices [144]. His approach seeks to make the original fixed point proof constructive, by finding approximate fixed points which correspond to approximately equilibrium prices for the market. His approach, while promising, had the disadvantage that it could take an exponentially long time before market prices were reached.

In 1976, Curtis Eaves managed to formulate the problem of finding these equilibrium prices as a linear complementarity problem [48]. Algorithms for computing solutions to these kinds of problems had already been established by that point, and so using, for example, Lemke's algorithm one could compute exact market prices. Again, though, this algorithm converges to equilibrium solutions only after an exponential amount of time in the worst case. Despite this, Eave's reformulation of the problem is still important today as it was the first proof that concretely demonstrated that if the inputs to the market are all rational numbers and utilities are linear, then the solution must be a rational number as well.

The first numerical success in computing equilibrium prices for perfectly competitive markets must be attributed to Irving Fisher and his famous 1891 Ph.D. thesis [26]. In his thesis, Fisher proposed his own framework for general equilibrium theory without any knowledge of Walras' work. His model, which turns out to be a special case of the more general model of Walras, seperates the market into buyers and sellers. Sellers are endowed with some set of goods and buyers are endowed with a certain amount of money which they bring to the market. Money is a special good for which sellers all have the same unit utility. We again seek a set of equilibrium market prices such that buyers and sellers can all maximize their utility without there being an excess of supply or demand.

What makes Fisher's research stand out from the work of his contemporaries is not the formulation of the theory of market equilibrium prices, though that is impressive in its own right. Fisher's unique accomplishment, was that he built a machine to actually compute these prices. The machine, a complex hydraulic device (see Figure 1–1 for an image from Fisher's thesis [26]) uses a system of cisterns and rods to encode the input endowments and express the output equilibrium prices. This approach to numerically computing equilibrium prices was decades ahead of his contemporaries and it would not be until the advent of rudimentary computers that Leontief was able to replicate a simplified version of this calculation in the 1930s [26].

Figure 1–1: Diagram of Fisher's Machine

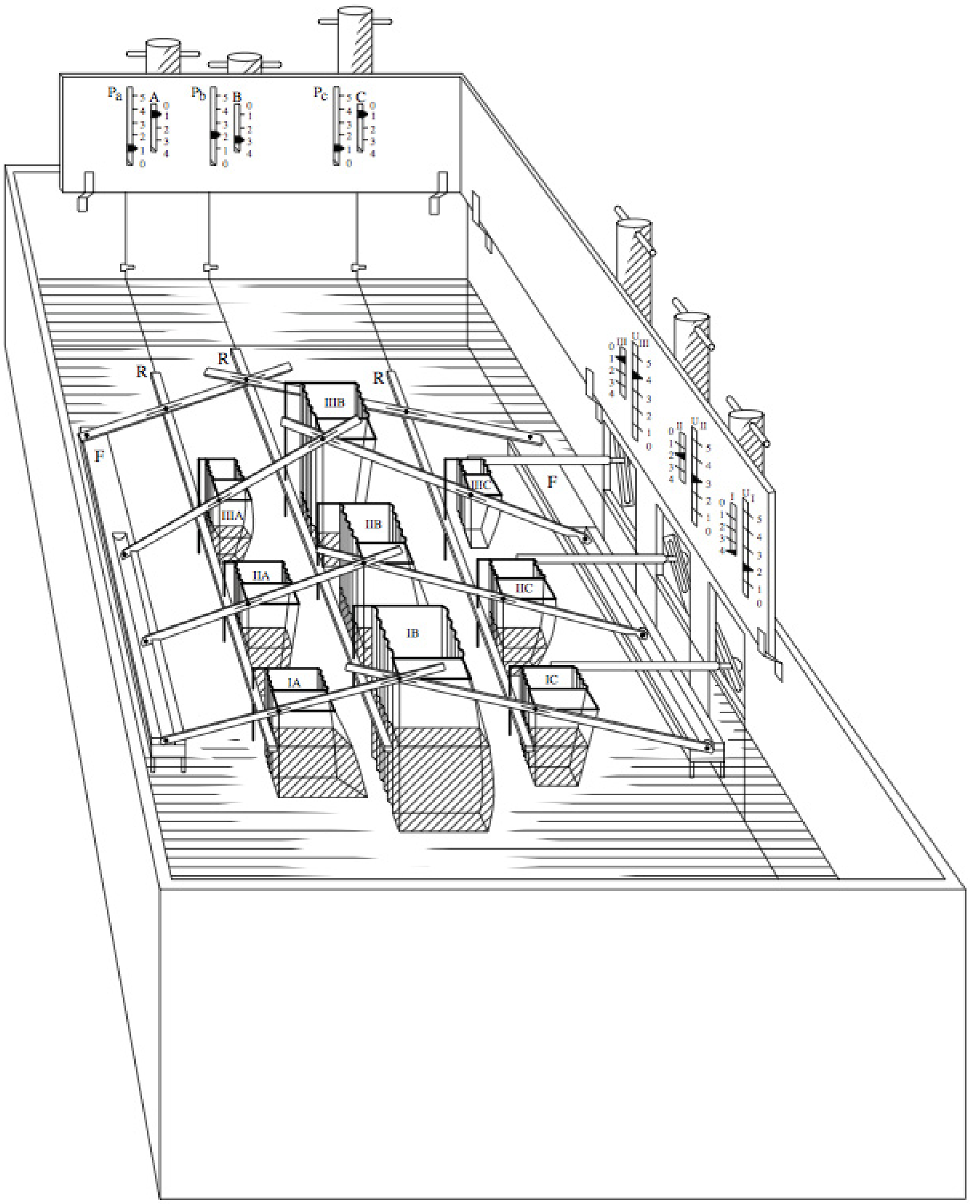

The Fisher model of general equilibrium theory proved to be much easier to analyze than Walras' model. In 1959, Eisenberg and Gale discovered a convex program whose optimal solutions corresponded exactly to market equilibrium prices in the case of linear utilities. A few years later, Eisenberg was able to extend this result to cover all homogenous utility functions. The existence of this convex program, coupled with Curtis Eave's proofs that all optimal solutions are rational, allows for a polynomial time algorithm to compute equilibrium market prices in the linear Fisher market model. This powerful result has led the Fisher model to become a key subject of study among 21st century algorithmic game theorists.[3]

In a 2001 paper [130], Christos Papadimitriou made a connection between market computation and the algorithmic theory of complexity when he presented a theorem that in the Arrow-Debreu model, if goods are integer valued instead of continuous, then it is NP-hard to determine if market equilibrium prices exist. However, a randomized polynomial time algorithm exists which can find ϵ-approximate prices (in expectation) to the equilibrium market prices. This algorithm makes use of a technique known as randomized rounding.

Algorithmic game theorists contend that the computational complexity of market problems is an important factor in determining the validity of equilibrium prices. Even if all actors in a market act as independent rational agents, the market is akin to a computer capable of parallel processing. By understanding both the complexity and approximability of market equilibria, we are able to discern if it is reasonable to

[3] Despite Fisher's success as an economist, he was not infallible. He is notorious for stating in a television interview that "Stock prices have reached what looks like a permanently high plateau" just days before the stock market crash of 1929.

expect that market agents will be able to reach an equilibrium by individual actions. Providing algorithms to find equilibria are often also important, as with the advent of the internet, many marketplaces have online components which these algorithms can help navigate.

In 2002, Devanur et al outline a combinatorial algorithm for finding equilibrium prices in the linear Fisher market model [47]. This algorithm extended the primal-dual methodology of linear progamming to the case of convex programming. By providing a combinatorial algorithm, rather than a mathematical programming result, Devanur et al's approach gives more insight into the structure of the market and provides an algorithm that is more flexible to changes in the problem setup. The insights of this algorithm inspire, for example, Nisan et al's work on Google's auctions for TV ads [126].

Other combinatorial algorithms for solving the linear Fisher market model have subsequently discovered. Garg and Kapoor outline an auction based algorithm [72] and Kelly and Vazirani demonstrate a technique that involves congestion analysis in a network [96]. These discoveries continue to shed light on the intricacies of market structures and also help us understand the connection that Fisher market mechanisms have to other areas of economics and computer science.

A convex programming formulation of the linear Arrow-Debreu model of general equilibrium theory was finally discovered in 2004 by Kamal Jain [87]. In this paper, Jain also demonstrates how this convex program can be solved using the ellipsoid method for the first polynomial time algorithm for finding market equilibrium prices, when they exist, in the more general Walrasian setting. Ye provides an intertior point algorithm using the same convex program in [179].

While we now have fast algorithms for market models with linear utilities, the context of more general utility functions is harder. For piecewise-linear concave utility functions, it has been demonstrated that finding market equilibrium prices is PPAD-complete ([36], [38], [173]). This complexity class, first defined by Christos Papadimitriou in 1994 [129], is believed to be intractable, leading algorithmic game theorists to question the validity of some of these market equilibria concepts. Much more work needs to be done to understand the true structure of even perfectly competitive markets and whether equilibrium prices can truly be found by market mechanisms.

CHAPTER 2
Economic Models

2.1 Models of Oligopoly

2.1.1 The Cournot Model

As mentioned in the Introduction, the original model of oligopolistic competition was formulated by Augustin Cournot in 1838 [43]. In this model, we assume a marketplace where a few sellers are selling identical, non-differentiated goods. Each seller chooses an amount of the good to produce and the price is set by market demand. We will define the model for the case of two sellers (a duopoly) but it is easily generalized to any number of market participants.

In this model, each player chooses some quantity of the good to produce, q_i, and pays some cost to produce it, $C_i(q_i)$. The price for the good is then set by the inverse demand function of the quantities produced by both players, $P(q_i + q_j)$. Each player i makes profit:

$$\Pi_i(q_i, q_j) = q_i P(q_i + q_j) - C_i(q_i). \tag{2.1}$$

For the remainder of this thesis, we will focus on the basic case of linear price and cost functions. In particular, we will only consider cost functions of the form $C_i(q_i) = c_i q_i + f_i$ for $q_i > 0$ where c_i is a constant marginal cost of production and f_i

is a fixed cost (this fixed cost is not paid if nothing is produced and so $C_i(0) = 0$). We will also only consider price functions of the form $P(q) = a - q$.[1]

The *Cournot Equilibrium* is a choice of quantites for each player such that each player is playing a best response to their opponent. In this game, a best response is a choice of quantity that maximizes a players profit given their opponent's strategy. In order to compute this equilibrium strategy, we begin by computing best responses.

Suppose player j chooses to produce quantity q_j^*. Player i's profit function is then

$$\Pi_i(q_i, q_j) = q_i(a - q_i - q_j) - c_i q_i - f_i. \tag{2.2}$$

Using first order conditions, we see that this is maximized when q_i is such that

$$a - 2q_i - q_j - c_i = 0 \tag{2.3}$$

i.e. when $q_i = \frac{a-q_j-c_i}{2}$. This equation describes the *best response curve* $BR_i(q_j)$ of player i. The equilibrium quantities are obtained when both players are playing best responses to each other, i.e. where these best response curves intersect (see Figure 2–1). Thus the Cournot Equilibrium in the linear duopoly case is (q_1^*, q_2^*) where:

$$q_i^* = \frac{a - 2c_i + c_j}{3}. \tag{2.4}$$

[1] More generally, linear price functions are of the form $P(q) = a - bq$, but we can assume $b = 1$ without loss of generality to simplify our calculations.

Figure 2–1: Best Response Curves in the Cournot Game

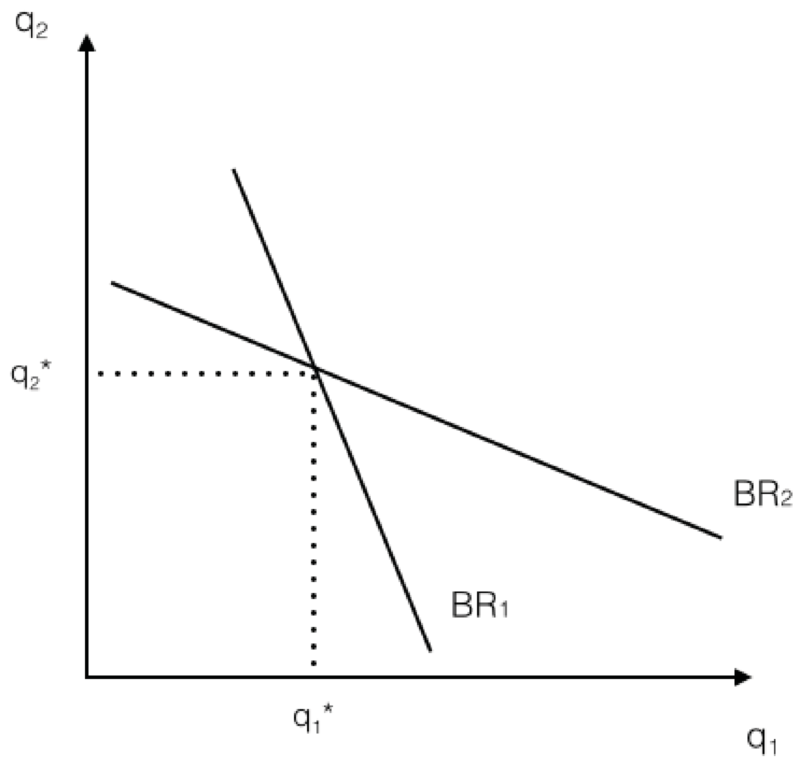

Their profit can be calculated to be

$$\Pi_i(q_i^*, q_j^*) = \left(\frac{a - 2c_i + c_j}{3}\right)^2 - f_i. \tag{2.5}$$

This equilibrium is unique in the absence of fixed costs. If there are positive fixed costs, then other equilibria may arise. For example, it is possible that player one's monopoly production would cause player two's maximum profit to fall under their fixed cost. In this instance, another equilibria would be the monopoly strategy for player one and for player two to produce nothing.

2.1.2 The Bertrand Model

The Bertrand model arose from a critique of the quantity strategization of Cournot[18]. Bertrand argues that is more natural and realistic to model oligopoly sellers as the price setter for goods in a market. As we did in the Cournot case, we will define the model for the duopolies and leave the generalization as clear.

Again, suppose we have two players each producing identical, non-differentiated goods. Player i has a cost function $C_i(q)$ based on the quantity he produces. We will again restrict our attention to the linear case so that

$$C_i(q) = \begin{cases} c_i q + f_i & \text{if } q > 0 \\ 0 & \text{otherwise} \end{cases} \tag{2.6}$$

Here, c_i is the marginal cost of production and f_i is a fixed cost which is paid only if production occurs at all. Each player then chooses the unit price p_i at which they will sell the good in the market. Since the goods are not differentiated, each consumer simply purchases the good from whomever charges the least. If both players charge the same price, then we assume that the market is shared evenly.

We assume that there is some function $D(p)$ that represents the market demand at a given price. As with the cost function, this thesis focusses on the case of linear demand so we assume $D(p) = a - p$.[2] This gives rise to the following profit function

[2] More generally a linear demand function is of the form $D(p) = a - bp$ but we may assume $b = 1$ without loss of generality to simplify our calculations.

for player i:

$$\Pi_i(p_i, p_j) = (p_i - c_i)D_i(p_i, p_j) - f_i \tag{2.7}$$

where $D_i(p_i, p_j)$ is the demand for player i's good under the current prices and is defined by

$$D_i(p_i, p_j) = \begin{cases} D(p_i) & \text{if } p_i < p_j \\ \frac{1}{2}D(p_i) & \text{if } p_i = p_j \\ 0 & \text{if } p_i > p_j \end{cases} \tag{2.8}$$

A natural consequence of this model is that there is only one Nash equilibrium and in it the player with the lower marginal cost gets the entire market by pricing at their opponent's marginal cost. If both players have the same marginal cost, then they will both price at cost and share the market. The profit of each player in this case will be $-f_i$.

This equilibrium result is sometimes called Bertrand's paradox. It implies that two oligopolists are enough to reduce pricing levels to that of perfect competition. This lack of pricing power of oligopolies seems counterintuitive to what is observed in real markets. Bertrand's model further implies that all such markets should, in fact, be monopolies as, if there is more than one participant, then almost all of the firms will be losing money.

There are a number of ways of enhancing the model to overcome this paradox. It is possible by introducing capacity constraints on firms, adding a time element to the game, or by allowing for product differentiation. For further discussion on each of these topics, see [168].

2.1.3 The Stackelberg Model

The Stackelberg model was formulated by Heinrich von Stackelberg in 1934 as an adaptation of the Cournot model of quantity competition [162]. It challenged the simultaneity assumption of Cournot's model and attempted to capture some of the dynamic elements of oligopoly competition.

In this model, the profit functions, price functions, and cost functions are identical to Cournot's model in Section 2.1.1. The Stackelberg model, however, separates the players into two types: leaders and followers. In the duopoly case, the model assumes that leader chooses their production quantity first and commits to it, after which followers make their choice with perfect information about the leader's choice.

Being able to commit first gives the leader an enormous advantage, as it forces the follower to optimize her profit on the leader's terms. From the follower's point of view, the optimization problem remains the same. Suppose player 1, the leader, produces q_1. Then the follower, player 2, in optimizing his profit is subject to the same first order condition (2.3). Thus he will choose $q_2 = \frac{a-q_1-c_2}{2}$.

However, the leader has a different problem than in the Cournot case. A rational leader would understand the optimization problem that the follower must face and could thus take it into account when he is choosing his strategy. Thus the leader must optimize the profit function:

$$\Pi_1(q_1) = \frac{q_1(a - q_1 + c_2)}{2} - c_1 q_1 - f_1. \tag{2.9}$$

His first order conditions is now:

$$\frac{a}{2} + \frac{c_2}{2} - q_1 - c_1 = 0. \tag{2.10}$$

The equilibrium strategy for the leader is to choose $q_1^* = \frac{a+c_2-2c_1}{2}$. In this case an optimizing follower must choose $q_2^* = \frac{a-3c_2+2c_1}{4}$. Notice that the leader captures a larger market share than in the Cournot game, at the expense of the follower. This is reflected in their equilibrium profits:

$$\Pi_1(q_1^*, q_2^*) = \left(\frac{a - 2c_1 + c_2}{2\sqrt{2}}\right)^2 - f_1 \tag{2.11}$$

$$\Pi_2(q_1^*, q_2^*) = \left(\frac{a - 3c_2 + 2c_1}{4}\right)^2 - f_2 \tag{2.12}$$

2.1.4 Comparison of Models

Let us compare the Cournot, Bertrand, and Stackelberg equilibrium outcomes to that of monopoly competition. We will consider the price, quantity produced, and total firm profits at equilibrium in a duopoly where sellers have equal marginal costs c and no fixed costs. We adapt the table below from [166].

Table 2–1: Comparison of Oligopoly Models

	Bertrand	Stackelberg	Cournot	Monopoly
Price	c	$(a+3c)/4$	$(a+2c)/3$	$(a+c)/2$
Quantity	$(a-c)$	$3(a-c)/4$	$2(a-c)/3$	$(a-c)/2$
Total Firm Profits	0	$3(a-c)^2/16$	$2(a-c)^2/9$	$(a-c)^2/4$

We assume $a > c$ for the demand (and inverse demand) function as, otherwise, the sellers' profits will always be negative and none will be incentivized to participate in the market.

Table 2–1 allows us to easily compare the quality of each equilibrium from the point of view of both buyers and sellers. Buyers naturally prefer Bertrand or Stackelberg competition, as more goods are produced at lower prices (indeed in this Bertrand example, twice as many goods are produced than at monopoly). Sellers would prefer to compete in Cournot or Monopoly settings to maximize their profits.

2.2 General Equilibrium Theory

General equilibrium theory is the study of market prices in contexts of perfect competition. The original models were formulated to discover if equilibrium prices existed that could simultaneously equate supply and demand throughout the market. In this section, we present the Fisher model of general equilibrium theory and discuss an important class of market utility functions.

2.2.1 The Fisher Market

A Fisher market $\mathcal{M}$, introduced by Irving Fisher in his 1891 PhD thesis, consists of a set $\mathcal{B}$ of buyers and and a set $\mathcal{G}$ of goods (owned by sellers). Let $n = |\mathcal{B}|$ and $g = |\mathcal{G}|$. Buyer i brings m_i units of money to the market and wants to buy a bundle of goods that maximizes her utility. Here, a non-decreasing, concave function $U_i : \mathbb{R}_+^g \to \mathbb{R}_+$ measures the utility she obtains from a bundle of goods. Without loss of generality, we may assume the aggregate quantity of each good on the market is one by scaling units appropriately.

Given prices $\boldsymbol{p} = (p_1, \ldots, p_g)$, where p_j is price of good j, each buyer demands a utility maximizing (an optimal) bundle that she can afford. The prices $\boldsymbol{p}$ are said to be a *market equilibrium* (ME) if agents can be assigned an optimal bundle such that demand equals supply, *i.e.* the market clears. Formally, let x_{ij} be the amount of good j assigned to buyer i. So $\boldsymbol{x}_i = (x_{i1}, \ldots, x_{ig})$ is her bundle. Then, $\boldsymbol{p}$ is an equilibrium price and $\boldsymbol{x}$ is an equilibrium allocation if the following two conditions are simultaneously satisfied.

1. `Supply = Demand`: $\forall j \in \mathcal{G},\ \sum_i x_{ij} = 1$ whenever $p_j > 0$.
2. `Utility Maximization`: $\boldsymbol{x}_i$ is a solution of $\max U_i(\boldsymbol{z})$ s.t $\sum_j p_j z_{ij} \leq m_i$.

We denote by y_{ij} the amount of money player i invests in item j after prices are set. Thus $y_{ij} = p_j x_{ij}$. Equivalently y_{ij} can be thought of as player i's demand for item j in monetary terms.

This model is a special case of the more general Arrow-Debreu model introduced in [7]. The more general model does not distinguish between buyers and sellers, instead proposing that all consumers to come to the market with a certain endowment of goods and are looking to trade goods so as to maximize their utility. It additionally allows for production to take place in the market via firms which take certain goods as inputs in order to produce others.

It was proved in [7] using fixed point techniques that as long as the utility functions of participants were continuous, quasi-concave, and satisfy non-satiation, a set of equilibrum market prices must exist in the Arrow-Debreu model.[3] Since the Fisher model is just a special case of Arrow-Debreu, equilibrium prices must exist in this model as well.

Theorem 1. *If the utility functions of buyers and sellers are continuous, quasi-concave, and satisfy non-satiation, then a set of market equilibrium prices exists.*

One of the advantages of the Fisher model over the Arrow-Debreu setting was that it proved to be computationally tractable. In [52], Eisenberg and Gale formulated a convex program that would capture equilibria prices and allocations in the Fisher model for linear utilities of the form $U_i(\boldsymbol{x}) = \sum_j u_{ij}x_{ij}$. This program is:

$$\begin{aligned} \max \quad & \textstyle\sum_i m_i \log U_i(\boldsymbol{x}) \\ \text{s.t.} \quad & \textstyle\sum_i x_{ij} \leq 1, \qquad \forall j \\ & x_{ij} \geq 0, \qquad \forall i, j. \end{aligned} \tag{2.13}$$

Any optimal allocation that solves this program is an equilibrium allocation. Equilibrium prices can be found as variables in the corresponding dual program. This program gives us a polynomial time algorithm for finding market prices (if they exist) in the linear Fisher game, via the ellipsoid method. Since 2002, a number of other methods have been formulated to compute equilibrium prices in this setting (see for example [47] or [72]).

[3] Non-satiation means that for every allocation of goods to player i, there is some other allocation that will increase his utility. Quasi-concavity means that if $u_i(x_i) > u_i(x'_i)$ then $u_i[tx_i + (1-t)x'_i] > u_i(x'_i)$ for any $0 < t < 1$.

2.2.2 Utility Functions

We are interested in analyzing a more general class of utility functions than just linear utilities. We focus on the important class of *Constant Elasticity of Substitution (CES)* utilities [161]. These functions have the form:

$$U_i(\boldsymbol{x}_i) = (\sum_j u_{ij} x_{ij}^{\rho})^{\frac{1}{\rho}} \tag{2.14}$$

for some fixed $\rho \leq 1$ and some coefficients $u_{ij} \geq 0$.

This class allows us to specify the elasticity of substitution for these markets as $\frac{1}{1-\rho}$. Hence, for $\rho = 1$, *i.e.* linear utilities, the goods are perfect substitutes; for $\rho \to -\infty$, the goods are perfect complements. As $\rho \to 0$, we obtain the well-known Cobb-Douglas utility function:

$$U_i(\boldsymbol{x}) = \prod_j x_{ij}^{u_{ij}} \tag{2.15}$$

where each $u_{ij} \geq 0$ and $\sum_j u_{ij} = 1$.

CES utility functions are widely used in economics as they are flexible enough to capture a large range of economic situations, but simple enough to remain computationally tractable. In fact, the Eisenberg-Gale program captures market equilibria prices for this more general class of functions. This was first proven by Eisenberg in [53][4] .

[4] In fact, Eisenberg showed that the program applied to all homogeneous utility functions, i.e. utilities which satisfied $U(\alpha\boldsymbol{x}) = \alpha U(\boldsymbol{x})$ for any scalar α.

A utility function is said to satisfy *weak gross substitutability* if increasing the price of one good cannot decrease demand for other goods. In a number of settings, utility functions having this property is the key to computational tractability. We have mentioned earlier that it is the condition under which the tâtonnement process converges for the Walras model ([6], [143]). Gross substitutability is also a critical condition in the theory of mechanism design. Tim Roughgarden states that "this condition captures the frontiers of tractability" for a wide range of auctions. We will find in our analysis in Chapter 5, that weak gross substituability is also the key condition for tractability in the Fisher Game. For CES utilities, this property is achieved when $0 < \rho \leq 1$ and Cobb-Douglas Utilities ($\rho \to 0$).

CHAPTER 3
Multimarket Price Wars

3.1 Introduction

In this chapter, we will explore the complexity of oligopolistic decision making. In particular, we are interested in the decision of when to start a price war or engage in predatory pricing in an oligoply markets. We focus on firms interacting in multiple markets (or a single segmentable market) as it allows us to model a broader and more realistic set of interactions.

A firm may initiate a price war by decreasing its prices in order to increase market share or to deter other firms from competing in particular markets. The firm suffers a short-term loss but may gain large future profits, particularly if the price war forces out the competition and allows it to price as a monopolist.

Price wars (and predatory pricing) have been studied extensively from both an economic and a legal perspective. A detailed examination of all aspects of price wars is far beyond the scope of this thesis. Rather, we focus on just one important aspect: the complexity of decision making in oligopolies (e.g. duopolies). Specifically, we consider the budget required by a firm in order to successfully launch a price war. This particular question is fundamental in determining the risk and benefits arising from predatory practices. Moreover, it arises naturally in the following two scenarios:

Entry Deterrence: How much of a war chest must a monopolist or cartel have on hand so that they are able to successfully repel a new entrant?

Competition Reduction: How much money must a firm or cartel have to force another firm out of business? For example, in a duopoly how much does a firm need to save before it can defeat the other to create a monopoly?

We formulate the War Chest Minimization Problem as a generalization of both of these scenarios and study the computational complexity of and approximation algorithms for this more general problem.

3.1.1 Background

Price wars and predatory pricing are tools that have been long associated with monopolies and cartels. The literature on these topics is vast and we touch upon just a small sample in this short background section.

Given the possible rewards for monopolies and cartels engaging in predatory behaviour, it is not surprising that it has been a recurrent theme over time. The late 19th century saw cartels engaging in predation in a plethora of industries. Prominent examples include the use of "fighting ships" by the British Shipping Conferences ([146], [133]) to control trade routes, the setting up of phoney independents by the American Tobacco Company to undercut smaller competitors [32]. Perhaps the most infamous instance, though, of a cartel concerns Standard Oil under the leadership of John D. Rockefeller ([112], [145], [44]). More recent examples of price wars include the cigarette industry [57], the airline industry [20], and the retail industry [25]. In the computer industry, Microsoft regularly faced accusations of predatory practices ([65], [102], [103]).

Antitrust legislation has been introduced in many countries to prevent anticompetitive behaviour like predatory pricing or oligopolistic collusion[1] . In the United States, the most important such legislation is the Sherman Act of 1890. One of the Act's earliest applications came in 1911 when the Supreme Court ordered the break-up of both Standard Oil and American Tobacco; more recently, it was applied when the Court ordered the break-up of American Telephone and Telegraph (AT&T) in 1982.[2]

Given that such major repercussions may arise, there is a need for a cloak of secrecy around any act of predation. This has meant the extent of predatory pricing is unknown and has been widely debated in the literature. Indeed, early economic work of McGee [112] suggested that predatory pricing was not rational. However, in Stigler's seminal work on oligopolies [165], price wars can be viewed as a break-down of a cartel, *albeit* they do not arise in equilibria because collusion can be enforced via punishment mechanisms. Moreover, recent models have shown how price wars can

[1] Whilst it is easy to see the negative aspect of cartels, it is interesting to note that there may even be some positive consequences. For example, it has been argued [64] that the predatory actions of cartels may *increase* consumer surplus.

[2] In 2000, a lower court also ordered the breakup of Microsoft for antitrust violations under the Sherman Act. On appeal, this punishment was removed under an agreed settlement in 2002.

be recurrent in a "functioning" cartel! For example, this can happen assuming the presence of imperfect monitoring [79] or of business cycles [138]. This is particularly interesting as recurrent price wars were traditionally seen as indicators of a healthy competitive market.[3]

Based primarily on the work of McGee, the US Supreme court now considers predatory pricing to be *generally implausible*.[4] As a result of this, and in an attempt to strike a balance between preventing anti-competitive behaviour and overly restricting normal competition, the Court applied the following strict definition to test for predatory practices.

1. The predator is pricing below its short-run costs.
2. The predator has a strong chance or recouping the losses incurred during the price-war.

The established way for the Court to test for the first requirement is the Areeda-Turner rule of 1975 [5] which established marginal cost (or, as an approximate surrogate, average variable cost) as the primary criteria for predatory pricing.[5] We will

[3] Therefore, should such behaviour also arise in practice it would pose intriguing questions for policy makers. Specifically, when is a price war indicative of competition and when is it indicative of the presence of a cartel or a predatory practice?

[4] See the 1986 case `Matsushita Electric Industrial Company vs Zenith Radio Corporation` and the 1993 case `Brooke Group Limited vs Brown and Williamson Tobacco Corporation`.

[5] We note that the Areeda-Turner rule may be inappropriate in high-tech industries because fixed costs there are typically high. Therefore, measures of variable

incorporate the Areeda-Turner rule as a legal element in our multimarket oligopoly models in Section 3.2.1. The second requirement essentially states that the "short-run loss is an investment in prospective monopoly profits" [56]. This requirement is typically simpler to test for in practice, and will be implicit in our models.

Finally, we remark that we are not aware of any other work concerning the complexity of price wars. One interesting related pricing strategy is that of loss-leaders which Balcan et al. [13] examine with respect to profit optimization. For the scale and type of problem we consider, however, using strategies that correspond to "loss-leaders" is illegal. Alternative models for oligopolistic competition and collusion in a single market setting can be found in the papers of Ericson and Pakes [58] and Weintraub et al. [177].

3.1.2 Our Results

A firm with price-making power belongs to an industry that is a monopoly or oligopoly. In Section 3.2, we develop three multimarket models of oligopolistic competition on top of the Bertrand, Cournot, and Stackelberg models introduced in Chapter 2. We then introduce the Minimum War Chest Problem to capture the essence of the Entry Deterrence and Competition Reduction scenarios outlined above.

---

costs may not be reflective of the presence of a price-wars. In fact, hi-tech industries may be particularly susceptible to predatory practices as large marginal profits are required to cover the high fixed costs. Consequently, predatory pricing can be used to inflict great damage on smaller firms.

In Section 3.3, we prove that this problem is NP-Hard in all three multimarket models under the legal constraints imposed by the Areeda-Turner rule. We emphasise that decision making is hard even under complete information. These hardness results utilise the fact that we have multiple markets. This assumption, however, is not essential. Decision-making can be hard in single-markets if either the number of firms is large or if the number of strategic options available to a firm is large. We give a simple example to illustrate this in Section 3.6.

We extend the hardness results in Section 3.4.2 to show that no multiplicative approximation guarantee can be obtained for the Minimum War Chest Problem, even in the simple case of linear cost, price, and demand functions. However, the situation for potential predators is less bleak than this result appears to imply. To see this we present two positive results in Section 3.4, assuming linear cost, price, and demand functions. First, the problem can be solved in polynomial time if the predator faces no fixed costs. In addition, for the Bertrand and Stackelberg models there is a natural way to separate the markets into two types, those where player one is making a profit and those in which she is truly fighting a price war. Our second result states that in these models, we can solve the problem on the former set of markets exactly and can find a fully polynomial time approximation scheme for the problem on the latter markets. This leads to a polynomial time algorithm with an arbitrarily small additive guarantee.

3.2 Models

3.2.1 Multimarket Models of Oligopoly

In this section, we formulate multimarket Bertrand, Cournot, and Stackelberg models as a generalization of the fundamental economic models introduced in Chapter 2. These allow for the investigation of more numerous and assorted interactions between firms.

A Multimarket Bertrand Model

Let us consider the following generalization of the asymmetric Bertrand model to multiple markets[6] . We will describe the model for the duopoly case, but again all of the definitions are easily generalizable. Suppose we have two players and n markets $m_1, m_2, ..., m_n$. Every player i has a budget B_i where a negative budget is thought of as the fixed cost for the firm to exist and a positive budget is thought of as a war chest available to that firm in the round. Every market m_k has a linear demand curve $D_k(p) = a_k - b_k p$ and each player i also has a marginal cost, c_{ik}, for producing one unit of good in market m_k. In addition, each player i has a fixed cost, f_{ik}, for each market m_k that she pays if and only if she enters the market, i.e. if she sets some finite price.

We model the price war as a game between the two players. A strategy for player i is a complete specification of prices in all the markets. Both players choose their

[6] We remark that this multimarket Bertrand model is also a generalization of the multiple market model used in the facility location game of Vetta [174].

strategies simultaneously. If $p_{ik} < \infty$ then we will say that player i enters market m_k. If player i chooses not to enter market m_k, this is signified by setting $p_{ik} = \infty$. The demand for each market then all goes to the player with the lowest price. If the players set the same price, then the demand is shared equally. Thus analogously to Section 2.1, if player i participates, then she gets profit Π_{ik} in market m_k where

$$\Pi_{ik}(p_{ik}, p_{jk}) = (p_{ik} - c_{ik})D_{ik}(p_{ik}, p_{jk}) - f_{ik} \tag{3.1}$$

and where D_{ik} is the demand for player i's good in market m_k and is defined as

$$D_{ik}(p_{ik}, p_{jk}) = \begin{cases} D_k(p_{ik}) & \text{if } p_{ik} < p_{jk} \\ \frac{1}{2}D_k(p_{ik}) & \text{if } p_{ik} = p_{jk} \\ 0 & \text{if } p_{ik} > p_{jk} \end{cases} \tag{3.2}$$

If player i chooses not to participate then her revenue and costs are both zero; thus, she gets 0 profit.

The sum of these profits over all markets is added to each player's budget. A player is eliminated if her budget is negative at the end of the round.

Multimarket Cournot and Stackelberg Models

We now formulate a multimarket version of the Cournot model. Again we will restrict ourselves to the case of the duopoly as the generalization is obvious. In this Cornout model, there are n independent Cournot markets $m_1, ..., m_n$. Each market m_k has a linear price function $P_k(q) = a_k - q$. Each player has a budget B_i, which serves the same role as in the Bertrand case. Each player i also has a cost function in every market $C_{ik}(q_{ik}) = c_{ik}q_{ik} + f_{ik}$, which consists of a marginal cost c_{ik} and a fixed cost f_{ik}.

As before we model the price war as a game. This time, a strategy for each player i is a choice of quantities q_{ik} for each market m_k. Again, both players choose a strategy simultaneously. We say that player i enters market m_k if $q_{ik} > 0$. As in the Bertrand case, a player pays the fixed cost f_{ik} if and only if they enter market m_k. Analogously to Section 2.1, player i then makes a profit in market m_k equal to

$$\Pi_{ik}(q_{ik}, q_{jk}) = q_{ik}P_k(q_{ik} + q_{jk}) - C_{ik}(q_{ik}). \tag{3.3}$$

Again, each player's aggregate profit is added to their budget at the end of the round. As above, a player is eliminated if her resulting budget is negative.

The multimarket Stackelberg model can then simply be adapted from the Cournot model. We define all of the quantities and functions as above. However, we now consider one player to be the leader and one to be the follower. The game is no longer simultaneous, as the leader gets to commit to a production level before the follower moves.

3.2.2 The War Chest Minimization Problem

We will examine the questions of entry deterrence and competition reduction in the two-firm setting. Thus, we focus on the computational problems facing (i) a monopolist fighting against a potential market entrant (entry deterrence) and (ii) a firm in a duopoly trying to force out the other firm (competition reduction). We model both these situations using the same duopolistic multimarket models of Section 3.2.1.

We remark that our focus on a firm rather than a cartel does not affect the fundamental computational aspects of the problem. This restriction, however, will allow us to avoid the distraction arising from the strategic complications that occur in ensuring coordination amongst members of a cartel.

Our game is then as follows. We assume that players one and two begin with budgets B_1 and B_2, respectively. They then play one of our three multimarket games. The goal of firm one is to stay/become a monopoly; if it succeeds it will subsequently be able to act monopolistically in each market. To achieve this goal the firm needs a non-negative payoff at the end of the game whilst its opponent has a negative payoff (taking into account their initial budgets). This gives us the following natural question:

War Chest Minimization Problem: *How large a budget B_1 does player one need to ensure that it can eliminate an opponent with a budget $B_2 < 0$.*

The players can play any strategy they wish *provided* it is legal, that is, they must abide by the Areeda-Turner Rule. All our results will be demonstrated under the assumptions of this rule, as it represents the current legal environment. However, similar complexity results can be obtained without assuming this rule.

Areeda-Turner Rule: *It is illegal for either player to price below their marginal cost in any market.*

Before presenting our results we make a few comments about the problem and what the legal constraints mean in our setting. First, notice that we specify a negative budget for player two but place no restriction on the budget for player one. This is natural for our models. We can view the budget as the money a firm initially has at its disposal minus the fixed costs required for it to operate; these fixed costs are additional to the separate fixed costs required to operate in any individual market. Consequently, if the second firm has a positive budget it cannot be eliminated from the game as it has sufficient resources to operate (cover its fixed costs) even without competing in any of the individual markets; thus we must constrain the second firm to have a negative budget. On the other hand, for the first firm no constraint is needed. Even if its initial budget is negative, it is plausible that it can still eliminate the second firm and end up with a positive budget at the end of the game, by making enough profit from the individual markets. Specifically, the legal constraints imposed by the Areeda-Turner rule may ensure that the second firm cannot maliciously bankrupt the first firm even if the first firm has a negative initial war-chest.

Second, since we are assuming that player one wishes to ensure success regardless of the strategy player two chooses, we will analyze the game as an asynchronous game where player two may see player one's choices before making her own. Player two will then first try to survive despite player one's choice of strategy. If she cannot do so, she will undercut player one in every market in an attempt to eliminate her also. To win the price war, player one must find strategies that keep herself safe and eliminate player two irrespective of how player two plays. Therefore, an optimal strategy for player one has maximum profit (i.e. minimum negative profit) amongst the collection of strategies that achieve these goals, assuming that player two plays maliciously.

Finally, the Areeda-Turner Rule has a straightforward interpretation in the Bertrand model of price competition, that is, neither player can set the price in any market below their marginal cost in that market! In models of quantity competition, however, the interpretation is necessarily less direct. For the Cournot model of quantity competition, we interpret the rule as saying that neither player can produce a quantity that will result in a price less than their marginal costly assuming the other player produces nothing, in other words $q_{ik} < a_k - c_{ik}$. This is the weakest interpretation possible for this simultaneous game. Finally, for the Stackelberg game, we assume that the restriction imposed by the Areeda-Turner rule is the same for player one as in the Cournot model, as she acts first and player two has not set a quantity when player one decides. Player two on the other hand, must produce a quantity so that her marginal price is greater than her marginal cost, given what player one has produced. In other words, for the Stackelberg game $q_{1k} < a_k - c_{1k}$ and $q_{2k} < a_k - q_{1k} - c_{2k}$.

3.3 Hardness Results

We are now in a position to show that the War Chest Minimization Problem is hard in all three models, even with linear price and demand functions. It follows that this problem is hard for more complex price and demand functions as well.

Theorem 2. *The War Chest Minimization Problem is NP-hard for the multimarket Bertrand model, even in the case with linear demand functions.*

Proof. We give a reduction from the knapsack problem. There we have n items, each with value v_i and weight w_i, and a bag which can hold weight at most W. In general, it is NP-hard to decide whether we can pack the items into the bag so that $\sum w_i \leq W$ and $\sum v_i > V$ for some constant V (where the sums are taken over packed items).

We will now create a multimarket Bertrand game based on the above instance. Suppose that there are n markets and each market m_k has the linear demand function

$$D_k(p) = 5\sqrt{v_k} - p. \tag{3.4}$$

Set player two's fixed costs to $f_{2k} = 0$ for all k and her marginal costs to $c_{2k} = 3\sqrt{v_k}$ for all k. Also set player one's marginal costs to $c_{1k} = 0$ for all k and her fixed costs to $f_{1k} = (25/4)v_k + w_k$ for all k. Set the budgets to be $B_1 = W$ and $B_2 = V - \sum_{k=1}^{n} v_k$.

We now calculate the monopoly prices for player one and player two. If player i wins market m_k at price p_{ik} then their profit in that market is

$$\Pi_{ik}(p_{ik}) = (p_{ik} - c_{ik})D_k(p_{ik}) - f_{ik} = -p_{ik}^2 + (c_{ik} + 5\sqrt{v_k})p_{ik} - 5\sqrt{v_k}c_{ik} - f_{ik}. \quad (3.5)$$

Taking derivatives, we see that the monopoly price for player i in market m_k is

$$p_{ik}^* = \frac{c_{ik} + 5\sqrt{v_k}}{2} \geq c_{ik}. \quad (3.6)$$

In particular, notice that the monopoly price for player one is

$$p_{1k}^* = \frac{5}{2}\sqrt{v_k} < 3\sqrt{v_k} = c_{2k}. \quad (3.7)$$

Moreover

$$p_{2k}^* = 4\sqrt{v_k} > 3\sqrt{v_k} = c_{2k}. \quad (3.8)$$

Thus, if player one enters market m_k then she can price at her monopoly price without fear that player two will undercut her. If she does not enter, then player two could price at her monopoly price to maximize revenue, as her fixed costs are zero. In the first case, player two earns 0 profit and player one earns monopoly profit

$$\begin{aligned} \Pi_{1k}(p_{1k}^*) &= -{p_{1k}^*}^2 + (c_{1k} + 5\sqrt{v_k})p_{1k}^* - 5\sqrt{v_k}c_{1k} - f_{1k} \\ &= -{p_{1k}^*}^2 + 5\sqrt{v_k}p_{1k}^* - f_{1k} \; = \; -w_k \end{aligned} \quad (3.9)$$

In the second case, player one earns zero profit while player two earns her monopoly profit

$$\begin{aligned}\Pi_{2k}(p^*_{2k}) &= -{p^*_{2k}}^2 + (c_{2k} + 5\sqrt{v_k})p^*_{2k} - 5\sqrt{v_k}c_{2k} - f_{2k} \\ &= p^*_{2k}(8\sqrt{v_k} - p^*_{2k}) - 5\sqrt{v_k}c_{2k} \;=\; v_k \end{aligned} \tag{3.10}$$

Thus, if player one could solve the War Chest Minimization Problem then she could determine whether or not there exists a set of indices K of markets that she should enter such that both of the following equations hold simultaneously:

$$W - \sum_{k \in K} w_k \;\geq\; 0 \tag{3.11}$$

$$V - \sum_{k=1}^{n} v_k + \sum_{k \notin K} v_k \;<\; 0 \tag{3.12}$$

Rearranging these equations, we obtain the conditions of the knapsack equations, namely $\sum_{k \in K} w_k \leq W$ and $\sum_{k \in K} v_k > V$. □

Theorem 3. *The War Chest Minimization Problem is NP-hard for the multimarket Cournot model, even in the case of linear price and cost functions.*

Proof. We again reduce from an instance of the knapsack problem. The Cournot game we create is as follows. Set $a_k = 6\sqrt{v_k}$, then for each market m_k let the price function be $P_k(q) = a_k - q$. We now set player one's marginal cost in market m_k to be $c_{1k} = 0$ and her fixed cost to be $f_{1k} = 4v_k + w_k$. Player two's marginal cost in market m_k is set to be $c_{2k} = 2a_k/3 = 4\sqrt{v_k}$ and her fixed cost is set to be $f_{2k} = 0$. Again, we set the budgets to be $B_1 = W$ and $B_2 = V - \sum_{k=1}^{n} v_k$.

Suppose now that player one has chosen which markets to enter and has, in particular, chosen to enter market m_k by producing quantity $q_{1k} > 0$. Consider player two's response. At first, player two will try to survive and will thus try to maximize her profit, given player one's quantity. She will consequently try to choose q_{2k} that maximizes $\Pi_{2k}(q_{1k}, q_{2k})$, call this quantity q_{2k}^+. By taking derivatives, we can calculate q_{2k}^+ to be $(a_k - 3q_{1k})/6$.

If player two calculates that she can't survive by choosing q_{2k}^+ in every market, then she will try to undercut player one in every market in an attempt to also drive her out. She will therefore choose $q_{2k} = q_{2k}^-$, the quantity which minimizes $\Pi_{1k}(q_{1k}, q_{2k})$. This can be achieved by making q_{2k} as large a possible; given the constraints of the Areeda-Turner rule this implies that $q_{2k}^- = a_k - c_{2k}$. Thus, we calculate $q_{2k}^- = a_k/3 = 2\sqrt{v_k}$.

Now, if we assume that player one enters market m_k (i.e. assume $q_{1k} > 0$) then by calculating the partial derivatives of $\Pi_{2k}(q_{1k}, q_{2k}^+)$ and $\Pi_{1k}(q_{1k}, q_{2k}^-)$ with respect to q_{1k}, we see that the quantity $q_{1k}^* = a_k/3 = 2\sqrt{v_k}$ minimizes the former and maximizes the latter. Therefore if player one chooses to enter market m_k she will produce quantity q_{1k}^*. So if player one enters market m_k then she, in the worst case, makes profit

$$\begin{aligned}\Pi_{1k}(q_{1k}^*, q_{2k}^-) &= q_{1k}^*(P_k(q_{1k}^* + q_{2k}^-) - c_{1k}) - f_{1k} \\ &= q_{1k}^*((6\sqrt{v_k} - 4\sqrt{v_k}) - 0) - f_{1k} \;=\; -w_k \end{aligned} \tag{3.13}$$

Against this, player two, in her best case, plays $q_{2k}^+ = (a_k - 3q_{1k}^*)/6 = 0$. This clearly gives her a profit $\Pi_{2k}(q_{1k}^*, q_{2k}^+) = 0$. On the other hand, if player one doesn't enter market m_k then she makes profit 0 in that market and player two makes her monopoly profit, which in this case is

$$\Pi_{2k}(0, q^*_{2k}) \;=\; q^*_{2k}(P_k(q^*_{2k}) - c_{2k}) - f_{2k} \;=\; q^*_{2k}(2\sqrt{v_k} - q^*_{2k}) - 0 \;=\; v_k \quad (3.14)$$

The proof follows. □

A similar proof holds for the Stackelberg case. We include the proof as it will be needed in Section 3.4.2.

Theorem 4. *The War Chest Minimization Problem is NP-hard for the multimarket Stackelberg model if player one is the Stackelberg leader, even in the case linear price and cost functions.*

Proof. We again reduce from the knapsack problem. Take any instance of the knapsack problem and define the quantities n, W, V, the w_is, and the v_is as in the proof of Theorem 2. We will now create a multimarket Stackelberg game based on the above instance. Set $a_k = 4\sqrt{v_k}$, and suppose that there are n markets and each market has price function $P_k(q) = a_k - q$. We now set player one's marginal cost in market m_k to be $c_{1k} = 0$ and her fixed cost to be $f_{1k} = 4v_k + w_k$. Player two's marginal cost in market m_k is set to be $c_{2k} = a_k/2 = 2\sqrt{v_k}$ and her fixed cost is set to be $f_{2k} = 0$. Finally, set the budgets to be $B_1 = W$ and $B_2 = V - \sum_{k=1}^{n} v_k$ as before.

Now consider the decision player one faces when deciding whether or not to enter market m_k. First notice that her monopoly quantity is $q^*_{1k} = a_k/2 = 2\sqrt{v_k}$ which we can calculate by maximizing $\Pi_{1k}(q_{1k}, 0)$ through simple calculus. Notice also that $a_k - q^*_{1k} - c_{2k} = 0$ and so, by the Areeda-Turner rule, player two cannot produce in any market in which player one is producing.

Thus, if player one enters any market then she will produce her monopoly quantity in that market and player two will not enter that market. In this case, player one makes profit

$$\begin{aligned} \Pi_{1k}(q^*_{1k}, 0) &= q^*_{1k}(P_k(q^*_{1k}) - c_{1k}) - f_{1k} \\ &= 2\sqrt{v_k}(2\sqrt{v_k} - 0) - (4v_k + w_k) \;=\; -w_k \end{aligned} \tag{3.15}$$

and player two makes profit $\Pi_{2k}(q^*_{1k}, 0) = 0$. On the other hand, if player one does not enter the market then player two will produce her monopoly quantity, $q^*_{2k} = a_k/4 = \sqrt{v_k}$, and will make profit

$$\begin{aligned} \Pi_{2k}(0, q^*_{2k}) &= q^*_{2k}(P_k(q^*_{2k}) - c_{2k}) - f_{2k} \\ &= \sqrt{v_k}(3\sqrt{v_k} - 2\sqrt{v_k}) - 0 \;=\; v_k \end{aligned} \tag{3.16}$$

Since player one did not enter, she will make profit 0. Thus we find ourselves back in the exact circumstances of the proof of Theorem 2. The rest of the proof follows. □

3.4 Algorithms

In this section, we explore algorithms for solving the War Chest Minimization Problem. We highlight a case where the problem can be solved exactly and explore the approximability of the problem in general. For the entirety of this section, we assume linear cost, demand, and price functions.

3.4.1 A Polynomial Time Algorithm in the Absence of Fixed Costs

All of the complexity proofs in Section 3.3 have a similar flavor. We essentially use the fixed costs in the markets to construct weights in a knapsack problem. In this section, we demonstrate that in the absence of fixed costs, it is computationally easy for a player to determine if they can win a multimarket price war even under the restrictions of the Areeda-Turner rule. This rule adds additional complications in this Stackelberg model, so we analyse that model first here. Again, we assume player one is the Stackelberg leader. Without fixed costs, the profit functions of both players in each market m_k are particularly simple:

$$\Pi_{ik}(q_{ik}, q_{jk}) = q_{ik}(a_k - q_{1k} - q_{2k} - c_{ik}) \tag{3.17}$$

As discussed, there are two strategies that player two may employ to prevent player one from winning the price war. She may play so as to survive or, if that is destined to fail, she may play so as to leave player one with a negative budget. In the former strategy, she will choose in every round and in every market the quantity, q_{2k}^{+}, that maximizes her own profit. In the latter strategy she will choose the quantity, q_{2k}^{-}, that minimizes player one's profit (while obeying the Areeda-Turner rule). By considering the partial derivatives of the players' profits, one can calculate q_{2k}^{+} and q_{2k}^{-} as functions of q_{1k}:

$$q_{2k}^{+} = \frac{a_k - q_{1k} - c_{2k}}{2} \quad (3.18)$$

$$q_{2k}^{-} = \begin{cases} a_k - q_{1k} - c_{2k} & \text{if } q_{1k} < a - c_{2k} \\ 0 & \text{otherwise} \end{cases} \quad (3.19)$$

The latter case for q_{2k}^{-} occurs if player one chooses a quantity so high that player two can choose nothing by the Areeda-Turner rule; this can only occur if $c_{1k} < c_{2k}$ as otherwise the Areeda-Turner rule itself prevents player one from choosing a sufficiently high quantity.

We now partition the markets into two sets: let $k \in A$ denote the set of markets for which $c_{1k} \leq c_{2k}$ and let $k \in B$ denote those markets where $c_{1k} > c_{2k}$. For the first subset A of markets, we will show that there is a natural choice of quantity for player one in every market. Namely, $q_{1k}^{+} = \max\{q_{1k}^{*}, a_k - c_{2k}\}$, where $q_{1k}^{*} = \frac{a_k - c_{1k}}{2}$ is player one's monopoly quantity. Clearly player one will never choose more than this as either (i) she is at her monopoly *and* player two can't enter or (ii) she is at a quantity that prevents player two from entering *and* increasing her quantity can only decrease her profit (since her profit is a concave quadratic). She will also never choose less than q_{1k}^{+} as she is either (i) at her monopoly quantity *and* preventing player two from entering or (ii) decreasing her quantity allows player two to enter the market with quantity $a_k - q_{1k} - c_{2k}$, resulting in player one selling fewer goods at a lower (or equal) price. Thus, in those markets A where player one is more competitive than player two, she will always enter at quantity q_{1k}^{+} and will always make a positive profit. Consequently, the optimal strategy for player one in these markets is clear. The problem, therefore, reduces to selecting quantities only in the subset B of markets where player one is less competitive.

So take a market $k \in B$. Then $q_{2k}^{-} = a_k - q_{1k} - c_{2k}$ always. Thus player one's profit, in the worst case is given by the linear function $q_{1k}(c_{2k} - c_{1k})$. So, again assuming that player two will first try to survive in every market and then try to undercut player one, the War Chest Minimization Problem for these markets is equivalent to the following quadratically constrained program:

$$\begin{array}{ll} \min & \sum_{k \in B} q_{1k}(c_{1k} - c_{2k}) \\ \text{s.t.} & \sum_{k \in B} (\frac{a_k - q_{1k} - c_{2k}}{2})^2 \leq -B_2 \\ & 0 \leq q_{1k} \leq a_k - c_{1k} \end{array}$$

We can solve this convex program in polynomial time.

The Cournot case is similar, this time with a convex quadratic objective function. In this case, q_{2k}^{+} is the same as above and the difference in Areed-Turner rule means that $q_{2k}^{-} = a_k - c_{2k}$. Thus, the War Chest Minimization Problem is equivalent to:

$$\begin{array}{ll} \min & \sum_{k} q_{1k}(q_{1k} + c_{1k} - c_{2k}) \\ \text{s.t.} & \sum_{k} (\frac{a_k - q_{1k} - c_{2k}}{2})^2 \leq -B_2 \\ & 0 \leq q_{1k} \leq a_k - c_{1k} \end{array}$$

The Bertrand case is even easier, as without fixed costs player one may enter every market and it is optimal for her to price at $p_{1k} = \max\{c_{1k}, \min\{p_{1k}^{*}, c_{2k} - \gamma\}\}$ where γ is some minimum increment of price and p_{1k}^{*} is player one's monopoly price in the market.[7] In other words, she will price either at her own marginal cost or just below player two's as prescribed by the Areeda-Turner rule, which provides a quick solution to the War Chest Minimization Problem. Thus, we have shown the following:

Theorem 5. *In the absence of fixed costs and assuming linear cost, price, and demand functions, the War Chest Minimization Problem in the Cournot, Bertrand, and Stackelberg models can be solved in polynomial time.*

3.4.2 An Inapproximability Result

In this section, we will explore approximation algorithms for the War Chest Minimization Problem. A first inspection is disheartening for would-be predators, as demonstrated by the following theorem.

Theorem 6. *It is NP-hard to obtain any (multiplicative) approximation algorithm for the War Chest Minimization Problem under the Bertrand, Stackelberg, and Cournot models.*

Proof. We prove this for the Stackelberg model - the other cases are similar. Let n, W, V, w_i, and v_i be an instance of the knapsack problem. Construct markets $m_1, ..., m_n$ exactly as in Theorem 4, with identical price functions, fixed costs, and marginal costs. Let W^* denote the optimal solution to the War Chest Minimization Problem in this case. Notice that $W^* > 0$ since all player one makes a negative profit in all of her markets. We now construct a new market m_{n+1} as follows. Let

[7] See the proof of Theorem 7 for a more detailed discussion.

$P_{n+1}(q) = 2\sqrt{W^*} - q$ be the price function. Let player one's fixed and marginal costs be $c_{1,n+1} = f_{1,n+1} = 0$. Let player two's marginal cost be $c_{2,n+1} = 2\sqrt{W^*}$ and let her fixed cost be an arbitrary nonnegative value. Then player one will clearly enter the market and produce her monopoly quantity, $q_{1,n+1} = \sqrt{W^*}$, thereby forcing player two to stay out of the market, by the Areeda-Turner rule. Thus player one will earn her monopoly quantity of W^* in this market. Consequently, the budget required for this War Chest Minimization Problem is zero. Any approximation algorithm would then have to solve this problem, and thereby the knapsack problem, exactly. □

3.4.3 Additive Approximation Guarantees

Observe that the difficulty in obtaining a multiplicative approximation guarantee arises due to conflict between markets that generate a loss for player one and markets that generate a profit. Essentially the strategic problem for player one is to partition the markets into two groups, α and β, and then conduct a price war in the markets in group α and try to gain revenue to fund this price war from markets in group β. This is still not sufficient because, in the presence of fixed costs, the optimal way to conduct a price war is not obvious even when the group α has been chosen. However, in this section we will show how to partition the markets and generate an arbitrarily small additive guarantee in the Bertrand and Stackelberg cases.

Given an optimal solution with optimal partition $\{\alpha^*, \beta^*\}$, let w_{α^*} be the absolute value of the sum of the profits of the markets with negative profit, and let w_{β^*} be the sum of the profits in positive profit markets. Then the optimal budget for player one is simply $OPT = w_{\alpha^*} - w_{\beta^*}$. For both the Bertrand and Stackelberg models, we will present algorithms that produce a budget of most $(1+\epsilon)w_{\alpha^*} - w_{\beta^*}$, for any constant ϵ. Observe this can be expressed as $OPT + \epsilon w_{\alpha^*}$, and since w_{α^*} represents the actual cost of the price war (which takes place in the markets in α^*), our solution is then at most OPT plus epsilon times the optimal cost of fighting the price war. Let's begin with the Bertrand model.

Theorem 7. *There is an algorithm that solves the War Chest Minimization Problem for the Bertrand model within an additive bound of ϵw_{α^*}, and runs in time polynomial in the input size and $\frac{1}{\epsilon}$, assuming linear demand functions.*

Proof. We begin by proving that we can find the optimal partition $\{\alpha^*, \beta^*\}$ of the markets. Towards this goal we show that there is a optimal pricing scheme for any market, should player one choose to enter the market. Using this scheme we will be able to see which markets are revenue generating for player one and which are not. This will turn out to be sufficient to obtain $\{\alpha^*, \beta^*\}$. This is because, in the Bertrand model, player two cannot make a profit in a market if player one does and vice versa and because player one needs a strategy that maintains a non-negative budget even if player two acts maliciously (but legally).

The pricing scheme for player one should she choose to enter market m_k is $p^+_{1k} = \max\{c_{1k}, \min\{p^*_{1k}, c_{2k} - \gamma\}\}$, where γ is the minimum increment of price and p^*_{1k} is player one's monopoly price. Certainly, she should not price below p^+_{1k} as either (i) it is illegal by the Areeda-Turner rule or (ii) she cannot increase her profit by doing so (as the profit function for player one is a concave quadratic in p_{1k}). She also should not price above p^+_{1k}. If she did then either (i) she cannot increase her profit (due to concavity) or (ii) player two could undercut her or increase her own existing profits in the market. Indeed, it is certain that player two will try to undercut her if player one succeeds in keeping player two's budget negative.

Given that we have the optimal pricing scheme for player one, we may calculate the profit she could make on entering a market assuming that player two acts maliciously. Let α be the set of markets where she makes a negative profit under these conditions, and let β be the set of markets where she makes a non-negative profit. Since all markets in β give player one a non-negative profit even if player two is malicious, she will clearly always enter all of them. Consequently, as we are in the Bertrand model, player two cannot make any profit from markets in β. Thus by entering every market in β player one will earn w_β profit, and this must be optimal for player one if the goal is to put player two out of business. So $\{\beta, \alpha\} = \{\beta^*, \alpha^*\}$ is an optimal partition.

It remains only to show that there is a fully polynomial time approximation scheme for the markets in α. We will prove this result by demonstrating an approximation preserving reduction of the War Chest Minimization Problem with only α-type Bertrand markets to the *Minimization Knapsack Problem*. Define w_k to be the negative of the profit earned by player one if she enters the market m_k and assuming player two undercuts if possible. By the above, she will price at $p_{1k} = p_{1k}^+$ and thus

$$w_k = \begin{cases} -(p_{1k}^+ - c_{1k})D(p_{1k}^+) + f_{1k} & \text{if } c_{1k} < c_{2k} \\ f_{1k} & \text{otherwise.} \end{cases} \tag{3.20}$$

Recall that w_k is non-negative for markets in α. Let p^*_{2k} be player two's monopoly price in market m_k and let Π^*_{2k} be her monopoly profit in that market. We also let $v_k = \Pi^*_{2k} - \Pi_{2k}(p^+_{1k})$, where $\Pi_{2k}(p^+_{1k})$ is the maximum profit that player two can achieve in market m_k if player one enters and prices at p^+_{1k}. The War Chest Minimization Problem is that of maximizing player one's profit (i.e. minimizing the negative of her profit) even if player two acts maliciously, while ensuring that player two's budget is always negative. So it can be expressed as

$$\begin{aligned}\min \quad & \textstyle\sum_k w_k y_k \\ \text{s.t.} \quad & B_2 + \textstyle\sum_k (\Pi^*_{2k}(1 - y_k) + \Pi_{2k}(p^+_{1k}) \cdot y_k) \leq 0 \\ & y_k \in \{0, 1\}\end{aligned}$$

Setting the constant C to be the sum of player two's budget and her monopoly profit in all of the markets, that is $C = B_2 + \sum_k \Pi^*_{2k}$, the problem can be rewritten as

$$\begin{aligned}\min \quad & \textstyle\sum_k w_k y_k \\ \text{s.t.} \quad & \textstyle\sum_k v_k y_k \geq C \\ & y_k \in \{0, 1\}\end{aligned}$$

Finally, since the w_k are non-negative, this formulation is exactly the minimization knapsack problem. The reduction is approximation preserving and so we are done as there is a fully polynomial time approximation scheme for the minimization knapsack problem [74]. □

We now turn to the Stackelberg problem.

Theorem 8. *There is an algorithm that solves the War Chest Minimization Problem for the Stackelberg model within an additive bound of ϵw_{α^*}, and runs in time polynomial in the input size and $\frac{1}{\epsilon}$, assuming linear cost and price functions.*

Proof. As we have seen, there are two strategies that player two may employ to prevent player one from winning the price war. She may play so as to survive or, if that is destined to fail, she may play so as to leave player one with a negative budget. As in Section 3.4.1, define the quantity q^+_{2k} to be the quantity that maximizes player two's own profit in every market and q^-_{2k} to be the quantity that minimizes player one's profit (while obeying the Areeda-Turner rule). As before, though now adjusting for fixed costs, we get:

$$q^+_{2k} = \begin{cases} \frac{a_k - q_{1k} - c_{2k}}{2} & \text{if } (\frac{a_k - q_{1k} - c_{2k}}{2})^2 \geq f_{2k} \\ 0 & \text{otherwise} \end{cases} \tag{3.21}$$

$$q^-_{2k} = \begin{cases} a_k - q_{1k} - c_{2k} & \text{if } q_{1k} < a - c_{2k} \\ 0 & \text{otherwise.} \end{cases} \tag{3.22}$$

We initially split the markets of the Stackelberg case into two sets: let $k \in A$ denote the set of markets for which $c_{1k} \leq c_{2k}$ and let $k \in B$ denote those markets where $c_{1k} > c_{2k}$. In the former case, if player one enters the market m_k then she will necessarily produce quantity $q^+_{1k} = \max\{q^*_{1k}, a_k - c_{2k}\}$, where $q^*_{1k} = \frac{a_k - c_{1k}}{2}$ is player one's monopoly quantity. The argument for this is identical to that in Section 3.4.1, as the fixed costs here do not change anything. Let β be the set of all markets for which player one's worst case profit, $\Pi_{1k}(q^+_{1k}, q^-_{2k})$, is now nonnegative. Clearly she will enter all of these markets, and again $\beta = \beta^*$. Let $\alpha = \alpha^*$ be the set of markets for which her worst case profit is negative. These include some of the markets in A and all of the markets in B, since player two as the follower can always force a price that is less than player one's marginal cost in these latter markets.

Again, player two will clearly enter each market in β^* and produce quantity q_{1k}^+, earning a positive profit of w_{β^*}. Thus, we need only find a fully polynomial time approximation scheme for the markets in α^*. So for the remainder of the proof, we will deal solely with the markets of α^*. By scaling, we may also assume that all of the variables and constants are integral.

As discussed above, the markets in $\alpha^* \cap A$ have a canonical choice of quantity for player one, q_{1k}^+. The worst case profit for player one in these markets will always be negative, by definition. Now let Π_{2k}^* be player two's monopoly profit in market m_k. Define V to be the sum of player two's monopoly profits in every market. Also define $v_k(q_{1k})$ to be the difference between player two's monopoly profit in market m_k and her maximum profit if player one enters the market with quantity q_{1k}. So, $v_k(q_{1k}) = \Pi_{2k}^* - \Pi_{2k}(q_{1k}, q_{2k}^+)$. Notice that $v_k(q_{1k})$ is monotonically nondecreasing in q_{1k}.

Define $w_k(q_{1k})$ to be player one's worst case cost (negative profit) if she chooses to produce quantity q_{1k} in market m_k. For those markets where $c_{1k} \leq c_{2k}$, there is a natural strategy for player one and so $w_k(q_{1k}) = -\Pi_{1k}(q_{1k}^+, q_{2k}^-) > 0$. For markets with $c_{2k} \leq c_{1k}$, we have $q_{2k}^- = a_k - q_{1k} - c_{2k}$ and so

$$w_k(q_{1k}) = \begin{cases} 0 & \text{if } q_{1k} = 0 \\ q_{1k}(c_{1k} - c_{2k}) + f_{1k} & \text{otherwise} \end{cases} \tag{3.23}$$

All of these weights are also non-negative.

The War Chest Minimization Problem requires player one minimize the cost of the markets she enters while keeping the sum of player two's budget and profits below zero. Since player two may "win" either by reducing player one's budget below zero or keeping her final budget nonnegative, player one needs to work with both her worst case costs and player two's best case profits. Thus the War Chest Minimization Problem, after some simple algebra, may be formulated as the problem of finding the integer vector $(q_{11}, q_{12}, ..., q_{1n})$ that solves

$$\begin{aligned} \min \textstyle\sum_k w_k(q_{1k}) & & \\ \text{s.t} \quad \textstyle\sum_k v_k(q_{1k}) & \geq V & \\ q_{1k} & \in \{0, q_{1k}^+\} & \text{if } c_{1k} < c_{2k} \\ 0 \leq q_{1k} & \leq a_k - c_{1k} & \text{if } c_{1k} \geq c_{2k}. \end{aligned}$$

The last constraint comes from the Areeda-Turner rule. We will refer to this problem as Stackelberg War Chest Minimization (SWCM). The weight of the vector $(q_{11}, q_{12}, ..., q_{1n})$ will mean $\sum_k w_k(q_{1k})$ and the value of the vector will mean $\sum_k v_k(q_{1k})$.

The remainder of this proof will be broken into parts. We first show that there is a pseudo-polynomial time dynamic program for SWCM. We then show how to use rounding techniques to obtain a polynomial time approximation scheme.

So let's describe the dynamic program. Let $\bar{W}$ be the maximum attainable weight. For each market m_i with $i \in \{1, ..., n\}$ and for each weight $w \in \{0, 1, ..., \bar{W}\}$, let $U_{i,w}$ denote the vector $(q_{11}, ..., q_{1n})$ such that $q_{1j} = 0$ for all $j > i$ which has total weight w and with the maximum value amongst all such vectors. Let $f(i,w)$ denote the value of $U_{i,w}$; if no such vector exists, then we set $f(i,w) = -\infty$. It is easy to calculate the base cases $f(1,w)$ for every w. We then get the recurrence:

$$f(i+1,w) = \max_{q_{1,i+1}} f(i, w - w_{i+1}(q_{1,i+1})) + v_i(q_{1,i+1}) \tag{3.24}$$

where the maximum is taken over the feasible values of $q_{1,i+1}$, where we understand that $f(n,w) = -\infty$ for all $w < 0$. Thus we get a dynamic program that solves SWCM exactly and whose running time is polynomial in $n, \bar{W}$, and $a_k - c_{1k}$ for those markets k with $c_{1k} \geq c_{2k}$.

This dynamic program is pseudo polynomial. We can make it polynomial by a suitable scheme to round the quantities and to round the weights. Rounding the quantities, we shall try to make the running time depend on $\log(a_k - c_{1k})$ instead of $a_k - c_{1k}$, for those markets k with $c_{1k} \geq c_{2k}$. To do this, we will restrict the possible feasible choices of quantity, in each of these market, in the following manner. First fix some $\delta_0 > 0$. For each interval $I = [0, a_k - c_{1k}]$, partition it into the subintervals $I_0 = \{0\}, I_1 = \{1\}, I_2 = (1,2], ..., I_i = (2^{i-2}, 2^{i-1}], ..., I_{\lceil \log(a_k - c_{1k}) + 1 \rceil} = (2^{\lceil \log(a_k - c_{1k}) - 1 \rceil}, a_k - c_{1k}]$. Each subinterval I_i, $i > 1$, is further partitioned into the minimum number of subintervals J_{ij} whose lengths are at most $\delta_0 2^{i-2}$. For each i, there are at most $\lceil \frac{1}{\delta_0} \rceil$ subintervals. For each quantity q_{1k} let $h_k(q_{1k})$ be the maximum value of the J_{ij} subinterval that contains q_{1k} (we define $h_k(0) = 0$ and $h_k(1) = 1$). Thus h_k maps the integer values of the interval $[0, a_k - c_{1k}]$ into a set of $O(\frac{1}{\delta_0} \log(a_k - c_{1k}))$ integers.

Now let $q = (q_{11}, ..., q_{1n})$ be any solution to SWCM. Since the objective function is linear, by replacing each q_{1k} with $h_k(q_{1k})$ we change the weight of the resulting vector by at most $\delta_0 w(q)$. By standard arguments, using these rounded quantities gives a $(1 + \delta_0)$ approximate algorithm whose running time is polynomial in $n, \frac{1}{\delta_0}, \bar{W}$ and $\log(a_k - c_{1k})$ for those markets k with $c_{1k} \geq c_{2k}$.

We can round the weights using a similar trick to obtain a $(1+\epsilon)$ approximation algorithm for SWCM whose running time is polynomial in $n, \log(a_1 - c_{11}), ..., \log(a_n - c_{1n}), \frac{1}{\epsilon}$ and $\log(\bar{W})$. This completes the proof. □

The approach taken here does not apply directly to the Cournot model. In particular, a more subtle rounding scheme is required there when player one is more competitive than player two. We conjecture, however, that a similar type of additive approximation guarantee is possible in the Cournot model.

3.5 Summary of Results

Here, we take a moment to summarize the complexity results we have achieved for the War Chest Minimization Problem:

1. All three models admit polynomial time algorithms in the absence of fixed costs.
2. The Bertrand and Stackelberg models are NP hard to solve within any multiplicative approximation, however an additive guarantee exists.
3. The Cournot model is NP hard to solve within any multiplicative approximation. The status of an additive guarantee is an open problem.

3.6 Single Market Case

Clearly, our hardness results require that there be a large number of markets (or submarkets). Whilst the multimarket problem is the most interesting one in our opinion, we remark that hardness results can be obtained even in the single-market case, provided that each firm has a sufficient number of strategic choices available to it. For example, in this section, we introduce a very simple modified single market model, where firms are able to invest in themselves by increasing their fixed cost to decrease their marginal cost. Despite the simplicity of this model, the War Chest Minimization Problem is trivially hard, indicating that more complex and realistic single market models will typically also be hard.

3.6.1 Hardness Result

Suppose player one and player two are competing in a single Bertrand market. Player two has a certain marginal cost c_2. Player one begins with a marginal cost c_1. However, she may choose to invest in any subset of n technologies each of which will cost her a fixed cost f_i but will reduce her marginal cost by λ_i. Suppose player one begins with a budget B_1 and may not spend more than this budget in technology investment. Player one wins the market from player two if she can reduce her marginal cost c_1 to below player two's c_2 within her budget constraints. This produces the problem:

Single Market War Chest Minimization Problem: *If the initial c_1 and c_2 are fixed, what is the minimum budget B_1 that player one needs so that she can win the* *market from player two?*

Theorem 9. *This problem is NP-hard but has a fully polynomial time approximation scheme.*

Proof. We prove the theorem by showing that this problem is completely equivalent to the minimization knapsack problem in an approximization preserving way. Notice that the problem can be formulated as

$$\begin{aligned} \min \quad & \textstyle\sum_i f_i x_i \\ \text{s.t.} \quad & c_1 - \textstyle\sum_i \lambda_i x_i < c_2 \\ & x_i \in \{0, 1\} \end{aligned}$$

But then if we write $v_i = f_i, w_i = \lambda_i, C = c_1 - c_2$ then we have reduced the problem to the minimization knapsack problem as seen in Section 3.4.3. This reduction clearly preserves approximation. □

CHAPTER 4
Lookahead Search

4.1 Introduction

In the previous chapter, we demonstrated how firms competing in oligopoly markets might face complex decisions. In this chapter, we explore how such complex economic games may be played in practice. To this end, we consider the strategy of lookahead search, described by Pearl [132] in his classical book on heuristic search as being used by "almost all game-playing programs". To understand the lookahead method and the reasons for its ubiquity in practice, consider an agent trying to decide upon a move in a game. Essentially, her task is to evaluate each of her possible moves (and then select the best one). Equivalently, if she knows the values of each child node in the game tree then she can calculate the value of the current node. However, the values of the child nodes may also be unknown! Recall two prominent ways to deal with this. Firstly, crude estimates based upon local information could be used to assign values to the children; this is the approach taken by *best response dynamics*. Secondly, the values of the children can be determined recursively by finding the values of the grandchildren. At its computational extreme, this latter approach in a finite game is Zermelo's algorithm - assign values to the leaf nodes of the game tree and apply backwards induction to find the value of the current node.[1]

Both these approaches are special cases of *lookahead search*: choose a local search tree T rooted at the current node in the game tree; valuations (or estimates thereof) are given to leaf nodes of T; valuations for internal tree nodes are then derived using the values of a node's immediate descendants via backwards induction; a move is then selected corresponding to the value assigned the root. For best response dynamics the search tree is simply the star graph consisting of the root node and its children. With unbounded computational power, the search tree becomes the complete (remaining) game tree used by Zermelo's algorithm.

In practice the actual shape of the search tree T is chosen *dynamically*. For example, if local information is sufficient to provide a reliable estimate for a current leaf node w then there is no need to grow T beyond w. If not, longer branches rooted at w need to be added to T. Thus, despite our description in terms of "backwards induction", lookahead search is a very forward looking procedure. Subject to our computational abilities, we search further forward only if we think it will help evaluate a game node. Indeed, in our opinion, it is this forward looking aspect that makes

[1] Often the values of the leaf nodes will be true values rather than estimates, for example when they correspond to end positions in a game.

lookahead search such a natural method, especially for humans and for dynamic (or repeated) games.[2]

Interestingly, the lookahead method was formally proposed as long ago as 1950 by Shannon [153], who considered it a practical way for machines to tackle complex problems that require "general principles, something of the nature of judgement, and considerable trial and error, rather than a strict, unalterable computing process". To illustrate the method, Shannon described in detail how it could be applied by a computer to play chess. The choice of chess as an example is not a surprise: as described the lookahead approach is particularly suited to game-playing. It should be emphasised again, however, that this approach is natural for all computationally constrained agents, not just for computers. Lookahead search is an instinctive strategic method utilised by human beings as well. For example, Shannon's work was in part inspired by De Groot's influential psychology thesis [80] on human chess players. De Groot found that all players (of whatever standard) used essentially the same thought process - one based upon a lookahead heuristic. Stronger players were better at evaluating positions and at deciding how to grow (prune or extend) the search tree but the underlying approach was always the same.

Despite its widespread application, there has been little theoretical examination of the consequences of decision making determined by the use of local search trees. The goal of this chapter is to begin such a theoretical analysis. Specifically, what are

[2] In contrast, strategies that are prescribed by axiomatic principles, equilibrium constraints, or notions of regret are much less natural for dynamic game players.

the quantitative outcomes and dynamics in various games when players use lookahead search?

4.1.1 Lookahead Search: The Model.

Having given an informal presentation, let's now formally describe the lookahead method. Here we consider games with sequential moves that have complete information. These assumptions will help simplify some of the underlying issues, but the lookahead approach can easily be applied to games without these properties.

We have a strategic game $G(\mathsf{P}, \mathcal{S}, \{\alpha_i : i \in \mathsf{P}\})$. Here P is the set of n players, S_i is the set of possible strategies for $i \in \mathsf{P}$, $\mathcal{S} = (S_1 \times S_2 \ldots \times S_n)$ is the strategy space, and $\alpha_i : \mathcal{S} \rightarrow R$ is the payoff function for player $i \in \mathsf{P}$. A *state* $\bar{s} = (s_1, s_2, \ldots, s_n)$ is a vector of strategies $s_i \in S_i$ for each player $i \in \mathsf{P}$.

Suppose player $i \in \mathsf{P}$ is about to decide upon a move. With lookahead search she wishes to assign a value to her current state node $\bar{s} \in \mathcal{S}$ that corresponds to the highest value of a child node. To do this she selects a search tree T_i over the set of states of the game rooted at $\bar{s}$. For each leaf node $\bar{l}$ in T_i, player i then assigns a valuation $\Pi_{j,\bar{l}} = \alpha_j(\bar{l})$ for each player j. Valuations for internal nodes in T_i are then calculated by induction as follows: if player p is destined to move at game node $\bar{v}$ then his valuation of the node is given by

$$\Pi_{p,\bar{v}} = \max_{\bar{u} \in \mathcal{C}(\bar{v})} [r_{p,\bar{v}} + \Pi_{p,\bar{u}}]. \tag{4.1}$$

Here, $\mathcal{C}(\bar{v})$ denotes the set of children of $\bar{v}$ in T_i, and $r_{p,\bar{v}}$ is some additional payoff received by player p at node $\bar{v}$. Should p choose the child $\bar{u}^* \in \mathcal{C}(\bar{v})$ then assume any non-moving player $j \neq p$ places a value of $\Pi_{j,\bar{v}} = r_{j,\bar{v}} + \Pi_{j,\bar{u}^*}$ on node $\bar{v}$. Then given values for children of the root node $\bar{s}$ of T_i, player i is thus able to compute the lookahead payoff $\Pi_{i,\bar{s}}$ which she uses to select a move to play at $\bar{s}$. (The method is defined in an analogous manner if players seek to minimise rather than maximise their "payoffs", e.g. minimize costs.)

After i has moved, suppose player j is then called upon to move. He applies the same procedure but on a local search tree T_j rooted at the new game node. Note that j's move may **not** be the move anticipated by i in her analysis. For example, suppose all the players use 2-lookahead search. Then player i calculates on the basis that player j will use a 1-lookahead search tree T'_j when he moves – because for computational purposes it is necessary that $T'_j \subseteq T_i$. But when he moves player j actually uses the 2-lookahead search tree T_j and this tree goes beyond the limits of T_i.

4.1.2 Lookahead Search: The Practicalities.

There is still a great deal of flexibility in how the players implement the model. For example

• **Dynamic Search Trees.** Recall that search trees may be constructed dynamically. Thus, the exact shape of the search tree utilized will be heavily influenced by the current game node, and the experience and learning abilities of the players. Whilst clearly important in determining gameplay and outcomes, these influences are a distraction from our focal point, namely, computation and dynamics in games in which players use lookahead search strategies. Therefore, we will simply assume here that each T_i is a breadth first search tree of depth k_i. Implicitly, k_i is dependent on the computational facilities of player i.

- **Evaluation Functions.** Different players may evaluate leaf nodes in different ways. To evaluate internal nodes, as described above, we make the standard assumption that they use a max (or min) function. This need not be the case. For example, a risk-averse player may give a higher value to a node (that it does not own) with many high value children than to a node with few high value children – we do not consider such players here.

- **Internal Rewards or Not: Path Model vs Leaf Model.** We distinguish between two broad classes of game that fit in this framework but are conceptually quite different. In the first category, payoffs are determined only by outcomes at the end of game. Valuations at leaf nodes in the local search trees are then just estimates of the what the final outcome will be if the game reaches that point. Clearly chess falls into this category. In the second category, payoffs can be accumulated over time - thus different paths with the same endpoints may give different payoffs to each player. Repeated games, such as industrial games over multiple time periods, can be modelled as a single game in this category. The first category is modelled by setting all internal rewards $r_{p,\bar{v}} = 0$. Thus what matters in decision making is simply the initial (estimated) valuations a player puts on the leaf nodes. We call this the *leaf (payoff) model* as an agent then strives to reach a leaf of T_i with as high a value as possible. The second category arises when the internal rewards, $r_{p,\bar{v}}$, can be non-zero. Each agent then wishes to traverse paths that allow for high rewards along the way.

More specifically, in this model, called the *path (payoff) model*, the internal reward is $r_{p,\bar{v}} = \alpha_p(\bar{v})$.

• **Order of Moves: Worst-Case vs Average-Case.** In multiplayer games, the order in which the players move may not be fixed. This adds additional complexity to the decision making process, as the local search tree will change depending upon the order in which players move. Here, we will examine two natural approaches a player may use in this situation: *worst case lookahead* and *average case lookahead*. In the former situation, when making a move, a risk-averse player will assume that the subsequent moves are made by different players chosen by an adversary to minimize that player's payoff. In the latter case, the player will assume that each subsequent move is made by a player chosen uniformly at random; we allow players to make consecutive moves. In both cases, to implement the method the player must perform calculations for multiple search trees. This is necessary to either find the worst-case or perform expectation calculations.

In practice, such versatility is a major strength and a key reason underlying the ubiquity of lookahead search in game-playing. For example, it accords well with Simon's belief, discussed in Section 4.1.4, that behaviours should be adaptable. For theoreticians, however, this versatility is problematic because it necessitates application-specific analyses. This will be apparent as we present our applications; we will examine what we consider to be the most natural implementation(s) of lookahead search for each game, but these implementations may vary each time!

4.1.3 Techniques and Results.

We want to understand the social quality of outcomes that arise when computationally-bounded agents use k-lookahead search to optimise their *expected* or *worst-case* payoff over the next k moves. Two natural ways we do this are via **equilibria** and via the study of **game dynamics**. In this thesis, we focus on the equilibria based approaches. To see an example of a dynamics based analysis refer to our full paper [116].

To explain the equilibrium approach, consider the following definition. Given a lookahead payoff function, $\Pi_{i,\bar{s}}$, a *lookahead best-response* move for player i, at a state $\bar{s} \in \mathcal{S}$, is a strategy s_i maximising her lookahead payoff, that is, $\forall s'_i \in S_i$: $\Pi_{i,\bar{s}} \geq \Pi_{i,(\bar{s}_{-i},s'_i)}$. (A move s'_i for player i, at a state $\bar{s} \in \mathcal{S}$, is *lookahead improving* if $\Pi_{i,\bar{s}} \leq \Pi_{i,(\bar{s}_{-i},s'_i)}$.) A *lookahead equilibrium* is then a collection of strategies such that each player is playing her lookahead best-response move for that collection of strategies. Our focus here is on pure strategies. Then, given a social value for each state, the *coordination ratio* (or price of anarchy) *of lookahead equilibria* is the worst possible ratio between the social value of a lookahead equilibrium and the optimal global social value.

It is worth noting that lookahead equilibria are generalizations of the more prevalent concept of a Nash equilibrium. By our definition, these concepts coincide for $k = 1$. Just as with Nash equilibria, there need not be a pure strategy lookahead equilibrium in any given game. We do not know of any results which address the question of the existence of a mixed strategy lookahead equilibrium.

The coordination ratio will allow us to discover when lookahead equilibria guarantee good social solutions, and how outcomes vary with different levels of foresight (k). In this thesis, we will focus on economic games of imperfect competition: particularly oligopoly games like the Cournot model and AdWord auctions. In our paper [116] we additionally analyze congestion games, valid-utility games, and a cost-sharing network design game.

We begin, in Section 4.2, with the Cournot duopoly game. Here two firms compete in producing a good consumed by a set of buyers via the choice of production quantities. We study equilibria in these simple games resulting from k-lookahead search. The equilibria for myopic game playing, $k = 1$, are well-understood in Cournot games. For $k > 1$, however, firms produce over 10% more than if they were competing myopically; this is better for society as it leads to around a 5% increase in social surplus. Surprisingly, the optimal level of foresight for society is $k = 2$. Furthermore, we show that Stackelberg behaviours arise as a special case of lookahead search where the firms have asymmetric computational abilities.

In Section 4.3, we examine strategic bidding in an AdWord generalised second-price auction, and study the social values of the allocations in the resulting equilibria. In particular, we show that 2-lookahead game playing results in the optimal outcome or a constant-factor approximate outcome, depending on the specifics of the modelf. This is in contrast to 1-lookahead (myopic) game playing which can result in arbitrarily poor equilibrium outcomes, and shows that more forward-thinking bidders would produce efficient outcomes.

Observe that our results show that lookahead search has different effects depending upon the game. It would be interesting to study further which game structures lead to more beneficial outcomes when longer foresight is used, and which game structures lead to more detrimental outcomes.

4.1.4 Background and Related Work.

This work is best viewed within the setting of *bounded rationality* pioneered by Herb Simon. In Rational Choice Theory a *rational* agent (or economic man) makes decisions via utility maximisation. Whilst the non-existence of economic man is not in doubt, rationality remains a central assumption in economic thought. This is typically justified using an *as if* as expounded by Friedman [68]: whether people are actually rational or not is unimportant provided their actions can be viewed in a way that is consistent with rational decision making - that is, provided agents act as if they are rational.[3] Friedman concluded that a model should be judged by its predictive value rather than by the realism of its assumptions. On this scale rationality often (but not always) does very well.

However, motivated by considerations of computational power and predictive ability, Simon [156] argued that "the task is to replace the global rationality of economic man with a kind of rational behaviour that is compatible with the access to

[3] For example, a consumer whose purchasing strategy allocates fixed proportions of her budget to specific goods (regardless of price levels) can be viewed as rational consumer with a Cobb-Douglas utility function!

information and the computational capacities that are actually possessed by organisms, including man, in the kinds of environments in which such organisms exist". He argued that, instead of optimising, agents apply heuristics in decision making. An example of this being the *satisficing* heuristic: agents search for feasible solutions, stopping when they discover an outcome that achieves an aspired level of satisfaction.[4] We remark that the use of a search phase provides a fundamental distinction between rational and boundedly rational agents. For rational agents the search is irrelevant as they will anyway make an optimal choice given the constraints of the problem. For agents of bounded rationality the form of the search can heavily influence decision making.

Interestingly, De Groot's work on chess players also heavily influenced Simon's general thinking on cognitive science.[5] This is exemplified in his famous book with Newell on human problem solving [125], where humans are viewed as information processing systems.

The label bounded rationality is currently used in a number of disparate areas some of which actually go against the main thrust of Simon's original ideas; see Selten [149] and Rubenstein [139] for some discussion on this point. Two schools of thought developed by psychologists, experimental economists, and behavioural

---

[4] Over time, and depending upon what is found in the search, this aspiration level may be changed.

[5] In fact, Simon sent his student George Baylor to help translate De Groot's work into English.

economists are, however, well worth mentioning here. First, the *Heuristics and Biases* program espoused by Kahneman and Tversky [170] and, second, the *Fast and Frugal Heuristics* program espoused by Gigerenzer [76]. Whilst both programs agree that humans routinely use simple heuristics in decision making, their philosophical outlooks are very different. The former program primarily looks for outcomes (caused by the use of heuristics) in violation of subjective excepted utility theory, and views such biases as a sign of irrationality likely to lead to poor decision making. In contrast, the latter program views the use of heuristics as natural and, in principle, entirely compatible with good decision making. For example, simple heuristics may be more robust to environmental changes and actually outperform methods based upon subjective excepted utility maximisation. As with the work of Simon, for the fast and frugal heuristics school, the actual quality of an heuristic is assumed to be dependent upon the search - how to search and when to stop searching - and the choice of decision rule after the search is terminated. Clearly, the lookahead heuristic can be viewed in this light: there is a search (via a local search tree), there is a "stopping rule" (determined, for example, by computational constraints and by the expertise of the player), and there is a decision rule (backwards induction).

The value of lookahead search in decision-making has been examined by the artificial intelligence community [123]; for examples in effective diagnostics and real-time planning see [98] and [147]. Lookahead search is also related to the sequential thinking framework in game theory [120, 163]. However, compared to these works and the research carried out by the two schools above, our focus is more theoretical and less experimental and psychological. Specifically, we desire quantitative performance guarantees for our heuristics.

Our research is also related to works on the price of anarchy in a game, and convergence of game dynamics to approximately optimal solutions [117, 78] and to sink equilibria [78, 60]. Numerous articles study the convergence rate of best-response dynamics to approximately optimal solutions [40, 62, 10, 22]. For example, polynomial-time bounds have been proven for the speed of convergence to approximately optimal solutions for approximate Nash dynamics in a large class of potential games [10], and for learning-based regret-minimisation dynamics for valid-utility games [22]. In another line of work, convergence of best-response dynamics to (approximate) equilibria and the complexity of game dynamics and sink equilibria have been studied [61, 2, 39, 159, 60, 115], but our thesis does not focus on these types of dynamics or convergence to equilibria.

Motivated by concerns of stability, convergence, and predictability of equilibria and game dynamics, various equilibrium concepts other than Nash equilibria have also been studied in the economics literature. Among them are correlated equilibria [8], stable equilibria [105], stochastic adjustment models [95], strategy subsets closed under rational behaviour (CURB set) [14], iterative elimination of dominated strategies, the set of undominated strategies, etc. Convergence and strategic stability of equilibria in evolutionary game theory is also an important subject of study. Many other game-theoretic models have been proposed to capture the self-interested behaviour of agents. As well as best-response dynamics, noisy best-response dynamics [54, 180, 118], where players occasionally make mistakes, simultaneous Nash dynamics [17], where all players change their strategies simultaneously, second-order Nash equilibria [19], where beginning with Nash equilibria the set of equilibria are recursively relaxed so that at any equilibrium there are no short, improving paths to worse equilibria, have all been studied.

In many other models the effect of learning algorithms [181] is examined, for example, regret minimisation dynamics [67, 82, 83, 23, 21, 22, 59] and fictitious play [30]. In most of these studies the most important factor is the stability of equilibria, and not measurements of the social value of equilibria. Furthermore, most of them are motivated by theoretical game theoretic concepts rather than practical game-playing, and none of the above works consider lookahead search.

4.2 Industrial Organisation: Cournot Competition

For our first example, we consider the classical Cournot model for duopolistic competition, which we introduced in Chapter 2. Our main result here is that the social surplus increases when firms are not myopic; surprisingly, social welfare is actually maximized when firms use 2-lookahead.

Recall that the Cournot model assumes players sell identical, nondifferentiated goods, and studies competition in terms of quantity (rather than price). Each player takes turns choosing some quantity of good to produce, q_i, and pays some marginal cost to produce it, c. In this chapter, we assume there are no fixed costs and that the marginal costs are symmetric to simplify the calculations. The analysis is easily extended to the non-symmetric case.

The price for the good is then set as a function of the quantities produced by both players, $P(q_i + q_j) = (a - q_i - q_j)$, for some constant $a > c$. On turn l, each player i makes profit: $\Pi_i^l(q_i, q_j) = q_i(a - q_i - q_j - c)$. In this form, the model then only has one equilibrium, called the *Cournot equilibrium*, where $q_i = (a - c)/3$ for each player. We may assume that $a = 1$ and $c = 0$. Then, at equilibrium, each player makes a profit of $\Pi^i(q_i, q_j) = q_i(1 - 2q_i)$. The consumer surplus is $2q_i^2$ and the social surplus (the sum of the firms profits and the consumer surplus) is $2q_i(1 - q_i)$.

4.2.1 Production under Lookahead Search.

We analyse this game when players apply k-lookahead search. In industrial settings it is natural to assume that payoffs are collected over time (as in a repeated game); thus, we focus upon the path model. We define this model inductively. In a k-step lookahead path model, each player i's utility is the sum of his utilities in the current turn and the $k - 1$ subsequent turns. He models the quantities chosen in the subsequent turns as though the player acting during those turns were playing the game with a smaller lookahead. More specifically, he assumes that the player acting in the t'th subsequent turn chooses their quantity to maximise their utility under a $k - t$ lookahead model. In order to rewrite this rigorously, let π_l^i be the contribution to his utility that player i expects on the lth subsequent turn (and π_0^i be the contribution to his utility that player i expects on his current turn), let π_l^j be the contribution to player j's utility that player i expects on the l'th subsequent turn, and let q_l^i (respectively, q_l^j) be the quantity that player i expects to choose (respectively, expects his opponent to choose) under this model.

Then in the path model, player i's expected utility function is $\Pi^i = \sum_{t=0}^{k-1} \pi_t^i$. Player j's expected utility function on player i's turn is $\Pi^j = \sum_{t=0}^{k-1} \pi_t^j$. Our aim now is to determine the quantities that player i expects to be chosen by both players in the subsequent turns and, thereby, determine the quantity he chooses this turn and the utility he expects to garner. To facilitate the discussion, it should be noted that unless noted otherwise, any reference to a "turn" refers to a turn during player i's calculation and not an actual game turn.

To simplify our analysis, we will define q_l to be the quantity chosen on turn l by whichever player is acting and Π_l to be the expected utility that player garners from turn l to turn k. So $\Pi_0 = \Pi^i$, $\Pi_1 = \sum_{t=1}^{k-1} \pi_t^j$, etc. We define $\overline{\Pi}_l$ to be the utility garnered from turn l to turn k by the player who does not act during turn l. So $\overline{\Pi}_0 = \Pi^j$, $\overline{\Pi}_1 = \sum_{t=1}^{k-1} \pi_t^i$, etc. It is clear that on each turn l, the active player is trying to maximise Π_l.

We are now ready to compute these quantities and utilities recursively. By our definition above, we have that $\Pi_k = q_k(1-q_k-q_{k-1})$ and $\overline{\Pi}_k = q_{k-1}(1-q_k-q_{k-1})$. Our definition also gives us the recursive formula for $l < k$ that $\Pi_l = q_l(1-q_l-q_{l-1})+\overline{\Pi}_{l+1}$ and $\overline{\Pi}_l = q_{l-1}(1-q_l-q_{l-1}) + \Pi_{l+1}$. Note that in each of these formulas, Π_l and $\overline{\Pi}_l$ are each functions of q_t for $t \geq l$; q_{l-1} is in fact fixed on the previous turn and is, therefore, not a variable in Π_l. It is now possible to calculate q_l recursively.

Lemma 1. *It holds that q_l is $\beta_l - \alpha_l q_{l-1}$, where $\beta_k = \alpha_k = \beta_{k-1} = \frac{1}{2}, \alpha_{k-1} = \frac{1}{3}$ and, for $l < k-1$,*

$$\beta_l = \frac{2 - \beta_{l+1} + \alpha_{l+1}\beta_{l+2} - \alpha_{l+1}\alpha_{l+2}\beta_{l+1}}{4 - 2\alpha_{l+1} - \alpha_{l+1}^2\alpha_{l+2}}, \quad \alpha_l = \frac{1}{4 - 2\alpha_{l+1} - \alpha_{l+1}^2\alpha_{l+2}} \tag{4.2}$$

Proof. We proceed by inducting down from q_k. Consider q_k which is the active player's choice on the final turn. As it is the final turn, he is acting myopically and so will choose q_k so as to maximise $\Pi_k = q_k(1 - q_k - q_{k-1})$. This parabola as a function of q_k is maximised when $q_k = \frac{1-q_{k-1}}{2}$. Doing a similar calculation for $\Pi_{k-1} = q_{k-1}(1 - q_{k-1} - q_{k-2}) + \overline{\Pi}_k$ gives us the desired values for β_{k-1} and α_{k-1}. We now assume the lemma for all $l > L$ and try to prove it for q_L. Recall the recursive formula $\Pi_L = q_L(1 - q_L - q_{L-1}) + \overline{\Pi}_{L+1}$. Taking the derivative of this with respect to q_L and setting it all equal to zero gives us

$$0 = (1 - 2q_L - q_{L-1}) + (1 - 2q_L - q_{L+1}) - \frac{\partial q_{L+1}}{\partial q_L} q_L - \frac{\partial q_{L+1}}{\partial q_L} q_{L+2} + \frac{\partial \Pi_{L+2}}{\partial q_{L+2}} \frac{\partial q_{L+2}}{\partial q_L}$$

The last term of the above sum is zero, since q_{L+2} is chosen so that $\frac{\partial \Pi_{L+2}}{\partial q_{L+2}} = 0$. Thus, if we plug in the inductive hypothesis into the above equation and simplify, we get

$$\begin{aligned} 2 - \beta_{L+1} + \alpha_{L+1}\beta_{L+2} - \alpha_{L+1}\beta_{L+2} &- \alpha_{L+1}\alpha_{L+2}\beta_{L+1} \\ &= (4 - 2\alpha_{L+1} - \alpha_{L+1}^2\alpha_{L+2})q_L + q_{L-1} \quad (4.3) \end{aligned}$$

This gives us the desired result. □

Our goal is now to calculate q_0 as this will tell us the quantity that player i actually chooses on his turn. From the above lemma, we can calculate q_0 if we can determine α_0 and β_0. Using numerical methods on the above recursive formula, we see that as $k \to \infty$, α_0 decreases towards a limit of $0.2955977\ldots$ and β_0 approaches a limit of $0.4790699\ldots$. These values also converge quite quickly; they both converge to within 0.0001 of the limiting value for $k \geq 10$. Thus, at a lookahead equilibrium, player i will choose $q_i \approx .0.4790699 - 0.2955977q_j$ and player j, symmetrically, will choose $q_j \approx 0.4790699 - 0.2955977q_i$. So each player will choose a quantity $q \approx 0.369767$. which is more than in the myopic equilibrium. Indeed, it is easy to show that for every $k \geq 2$, each player will produce more than the myopic equilibrium.

This is illustrated in Figure 4–1. Observe the quantity produced does not change monotonically with the length of foresight k, but it does increase significantly if non-myopic lookahead is applied at all. Consequently, in the path model looking ahead is better for society overall but worse for each individual firm's profitability (as the increase in sales is outweighed by the consequent reduction in price).

Figure 4–1: How output varies with foresight k

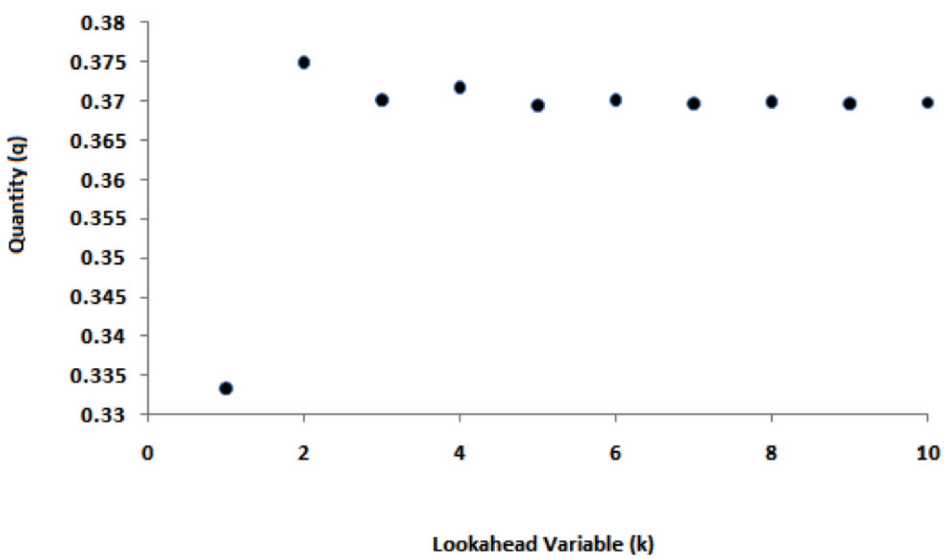

Theorem 10. *For Cournot games under the path model, output at a k-lookahead equilibrium peaks at $k = 2$ with output 12.5% larger than at a myopic equilibrium ($k = 1$). As foresight increases, output is 10.9% larger in the limit. The associated rises in social surplus are 5.5% and 4.9%, respectively,*

4.2.2 Stackelberg Behaviour.

We could also analyse this game under the leaf model, but this model is both less realistic here and trivial to analyse. However, it is interesting to note that for the leaf model with asymmetric lookahead, where player i has 2-lookahead and player j has 1-lookahead, we get the same equilibrium as the classic Stackelberg model for competition. Thus, the use of lookahead search can generate leader-follower behaviours.

4.3 Generalised Second-Price Auctions

For our second example, we apply the lookahead model to generalised second-price (GSP) auctions. Our main results are that outcomes are provably good when agents use additional foresight; in contrast, myopic behaviour can produce very poor outcomes.

The auction set-up is as follows. There are T slots for sale with click-through rates $c_1 > c_2 > ... > c_T > 0$, that is, higher indexed slots have lower click-through rates. There are $n > T$ players bidding for these slots, each with a private valuation v^i. Each player i makes a bid b^i. Slots are then allocated via a *generalised second price auction*. Denote the jth highest bid in the descending bid sequence by b_j, with corresponding valuation v_j. The jth best slot, for $j \leq T$, is assigned to the jth highest bidder who is charged a price equal to b_{j+1}. The T highest bidders are called the "winners". According to the pricing mechanism, if bidder i were to get slot t in the final assignment, then he would get utility $u^i_t = (v^i - b_{t+1})c_t$. We denote a player i's utility if he bids b^i by $u^i(b^i)$ (the other players bids are implicit inputs for u^i).

This auction is used in the context of keyword ad auctions (e.g, Google AdWords) for sponsored search. Given the continuous nature of bids in the GSP auction, the best response of each bidder i for any vector of bids by other bidders corresponds to a range of bid values that will result in the same outcome from i's perspective. Among these set of bid values, we focus on a specific bid value b^i, called the *balanced* bid [34]. The balanced bid b^i is a best-response bid that is as high as possible such that player i cannot be harmed by a player with a better slot undercutting him, i.e. bidding just below him. It is easy to calculate that for player i in slot t, $1 \leq t < T$, the only balanced bid is

$$b^i = (1 - \frac{c_t}{c_{t-1}})v^i + \frac{c_t}{c_{t-1}} b_{t+1}. \tag{4.4}$$

An important property of balanced bidding is that each "losing" player i (one not assigned a slot) should bid truthfully, that is $b^i = v^i$. To see this add dummy slots with $c_t = 0$ if $t > T$. The player who wins the top slot should also bid truthfully under balanced bidding. Balanced bidding is the most commonly used bidding strategy [34, 111]. For some intuition behind this, note that balanced bidding has several desirable properties. For a competitive firm, bidding high obviously increases the chance of obtaining a good slot. Within a slot this also has the benefit of pushing up the price a competitor pays without affecting the price paid by the firm. On the other hand, bidding high increases the upper bound on the price the firm may pay, leading to the possibility that the firm may end up paying a high price for one of the less desirable slots. Balanced bidding eliminates the possibility that a change in bid from a higher bidder can hurt the firm. (Clearly, it is impossible to obtain such a guarantee with respect to a lower bidder.) Thus, balanced bidding provides some of the benefits of high bidding at less risk. Balanced bidding naturally converges to Nash equilibria unlike other bidding strategies such as altruistic bidding or competitor busting [34]. Moreover, the other bidding strategies would require some discretization of players' strategy space in order to analyse the best response dynamics [34, 111]. Consequently, balanced bidding is the most natural strategy choice for our analysis.

For this auction problem, we consider only the leaf model. The leaf model seems more natural than the path model for a single auction as players are interested in the final allocation output by the auction (there are no intermediary payoffs). We analyse both worst-case and average-case lookahead; depending upon the level of risk-aversion of the agents both cases seem natural in auction settings.

Let player i's lookahead payoff (or utility) at bid b^i with respect to player j, denoted by $u^{ij}(b^i)$, be player i's payoff (or utility) after player j makes a best-response move. In the worst-case lookahead model, we define player i's lookahead payoff for a vector $\bar{b}$ of bids as $\Pi_{i,\bar{b}} = \tilde{u}^i(b^i) = \min_j u^{ij}(b^i)$. In the average-case lookahead model, player i's lookahead payoff $\Pi_{i,\bar{b}}$ for a bid vector $\bar{b}$ is $\Pi_{i,\bar{b}} = \bar{u}^i(b^i) = \frac{1}{n}\sum_j u^{ij}(b^i)$. Changing strategy from bid b^i to bid $\bar{b}^i$ is a *lookahead improving* move if lookahead utility increases, i.e., $\bar{u}^i(\bar{b}^i) > \bar{u}^i(b^i)$. We are at a *lookahead equilibrium* if no player has a lookahead improving move.

It is known that the social welfare of Nash equilibria for myopic game playing can be arbitrarily bad [34] unless we disallow over-bidding [108]. Here, we prove the advantage of additional foresight by showing that 2-lookahead equilibria have much better social welfare. In particular, we show that all such equilibria are optimal in the worst-case lookahead model, and all such equilibria are constant-factor approximate solutions in the average-case lookahead model.

4.3.1 Worst-Case Lookahead.

Our proof for the worst-case lookahead model can be seen as a generalisation of the proof of [31] for a slightly different model. We start by proving a useful lemma in this context.

Lemma 2. *Label the players so that player i is in slot i (so $v_i = v^i$ for all i), and suppose there is a player t such that $v^t < v^{t+1}$. Then player t myopically prefers slot $t+1$ to slot t.*

Proof. Suppose not. Then, as player t does not myopically prefer slot $t+1$ we have

$$(v_t - b_{t+1})c_t \geq (v_t - b_{t+2})c_{t+1} \tag{4.5}$$

By definition, $b_{t+1} = v_{t+1} - \frac{c_{t+1}}{c_t}(v_{t+1} - b_{t+2})$. Plugging this in gives

$$\begin{aligned}(v_t - b_{t+2})c_{t+1} &\leq \left(v_t - \frac{c_t - c_{t+1}}{c_t}v_{t+1} - \frac{c_{t+1}}{c_t}b_{t+2}\right)c_t \\ &< \left(\frac{c_{t+1}}{c_t}v_t - \frac{c_{t+1}}{c_t}b_{t+2}\right)c_t \\ &= (v_t - b_{t+2})c_{t+1}\end{aligned} \tag{4.6}$$

Thus we obtain our desired contradiction. Note that the strict inequality above follows directly from the fact that $v^t < v^{t+1}$. □

An equilibrium is *output truthful* if the slots are assigned to the same bidders as they would be if bidders were to bid truthfully. It is easy to verify that an allocation optimizes social welfare if and only if it is output truthful. Thus to prove 2-lookahead equilibria are socially optimal it suffices to show they are output truthful.

Theorem 11. *For GSP auctions, any* 2*-lookahead equilibrium gives optimal social welfare in the worst-case, leaf model.*

Proof. We proceed by contradiction. Consider a non-output-truthful 2-lookahead equilibrium. Again, label the players so that the player i is in slot i. Amongst all the winning players, take the one with the lowest valuation, v_i. First suppose that v_i is not amongst the T highest valuations. Then, there is a losing player with a higher value than v_i. But this player is bidding his value, as a result of balanced bidding. Consequently, player i's utility must be negative, a contradiction.

Thus, we may assume that v_i is amongst the T highest valuations; specifically it must have exactly the Tth highest valuation. We will show that player i moving into slot T is a lookahead improving move. Notice that the lookahead value for player i staying in slot i is at most the myopic value of staying in that slot. This follows from the fact that the myopic play of a losing player cannot improve the utiility of player i. Hence, it suffices to show that the lookahead value of changing slots is better than the myopic value of staying in slot i.

By several applications of Lemma 2, we see that player i myopically prefers slot T to slot i. However, in moving to slot T, player i will still make a balanced bid. Thus, no other winning player may reduce i's utility by undercutting him. Also, no losing player j wants to move to a winning slot as they can only be left with negative utility - since j cannot then be amongst the T highest valuations. So moving to slot T is a lookahead improving move for player i.

If player i were originally in slot T, then the entire argument can be applied with regards to slots 1 to $T - 1$. Inductively, we then conclude that in any non-output-truthful equilibrium, there is a lookahead improving move, which is a contradiction. This gives us the desired result. □

4.3.2 Average Case Lookahead.

Next, we consider the average-case lookahead model and show that the above theorem does not hold for this case.

Theorem 12. *In GSP auctions, there exist* 2*-lookahead equilibria that are not output-truthful in the average-case, leaf model.*

Proof. Consider the following example with $n = T = 4$. Let the click-through rates be $c_1 = 35, c_2 = 26, c_3 = 25$, and $c_4 = 20$. Let the valuations be $v_1 = 82, v_2 = 83, v_3 = 100, v_4 = 93$. Starting with the highest slot and working to the lowest, let bidder i bid the balanced bid for slot i. It can be verified that this turns out to be a non-output-truthful equilibrium. □

Despite this negative result, 2-lookahead equilibria cannot have arbitrarily bad social welfare.

Theorem 13. *In GSP auctions, the coordination ratio of* 2*-lookahead equilibria is constant in the average-case, leaf model.*

Proof. Suppose that we are at an equilibrium. Let v_{i^*} be the i^{th} highest valuation, let player i^* denote the corresponding player, let b_{i^*} denote their bid, and c_{i^*} be the click through rate of the slot they currently occupy. We recall that v_i denotes the player in slot i and it has click through rate c_i and bid b_i. The social welfare of a set A of players is $\sum_{i \in A} v_i c_i$. Thus, by the above definitions, the optimal social welfare is $\sum_i v_{i^*} c_i$.

Now, choose $\alpha, \beta < 1$ such that $(1 - \alpha)^2 > m\beta$ for some m to be chosen later. Let I be the set of indices i that satisfy both $v_i < \alpha v_{i^*}$ and $c_{i^*} < \beta c_i$. Note that for all $i \notin I$ the pair of players $\{v_i, v_{i^*}\}$ contribute at least $\min\{\alpha, \beta\} v_{i^*} c_i$ to OPT. So if I is empty, then we have achieved a constant coordination ratio. We may thus suppose I is not empty and choose $i \in I$.

Consider c_{i^*-1}. As we assume "balanced" bidding, $b_{i^*} \geq (1 - \frac{c_{i^*}}{c_{i^*-1}}) v_{i^*}$. Since $b_{i^*} < b_i < v_i < \alpha v_{i^*}$ by assumption, we have $c_{i^*-1} < \frac{1}{1-\alpha} c_{i^*}$. Choose $m > 1$. We first prove the following claim.

Claim 1. *For all $i \in I$, we have $c_{i+1} \leq \frac{c_i}{m}$.*

Proof. Suppose $c_{i+1} > \frac{c_i}{m}$, for some $i \in I$. We will show that player i^* moving into slot i is then lookahead improving. Consider his lookahead utility for staying put. Ignoring a repeat move for player i^*, which occurs with probability $\frac{1}{n}$, player i^*'s utility in every other circumstance is at most $c_{i^*-1} v_{i^*}$, as other players can improve his position by at most one. On the other hand, if player i^* moves into slot i then his lookahead utility is at least $c_{i+1}(v_{i^*} - b_i)$; he wins at least slot $i + 1$ and pays at most his bid. If player i is chosen to repeat his move then his utility is the same for both cases (as he will then simply play a best response move). Thus, it is enough for us to show that

$$c_{i+1}(v_{i^*} - b_i) > c_{i^*-1}v_{i^*} \tag{4.7}$$

However $b_i < v_i < \alpha v_{i^*}$ and putting this together with the above inequalities gives

$$c_{i+1}(v_{i^*} - b_i) > \frac{c_i}{m}(1-\alpha)v_{i^*} \geq \frac{\beta}{1-\alpha}c_i v_{i^*} > \frac{1}{1-\alpha}c_{i^*}v_{i^*} > c_{i^*-1}v_{i^*} \tag{4.8}$$

We are now done, by our choice of α and β, and have shown that player i^* moving into slot i is a lookahead improving move. This contradicts the fact we are at an equilibrium. $\square$

Thus we have established that for all $i \in I$, $c_{i+1} < \frac{c_i}{m}$. Thus, we can bound the optimal social welfare contributed by the slots $i \in I$ by $\frac{m}{m-1}c_{i_0}v_{i_0*}$ where $i_0 = \min_{i \in I} i$.

Now if $1 \notin I$ then we have achieved our constant coordination ratio since then either $c_1 v_1 > \alpha c_1 v_{1^*}$ or $c_{1^*}v_{1^*} \geq \beta c_1 v_{1^*}$. Hence, we are guaranteed at least $\min\{\alpha, \beta\}c_1 v_{1^*} \geq \min\{\alpha, \beta\}c_{i_0}v_{i_0*}$, that is, a least a constant factor of the social welfare from all the slots in I in the optimal allocation. So we suppose $1 \in I$.

Choose $\alpha_1 = \frac{m}{m-1}\alpha$ and consider the player currently in slot 2. By this choice of α_1, we ensure that this player does not have value more than $\alpha_1 v_{1^*}$. To see this, recall the player is bidding in a balanced manner and so, by Claim 1, his bid b_2 satisfies

$$v_2 \geq b_2 \geq (1 - \frac{c_2}{c_1})v_2 \geq (1 - \frac{1}{m})v_2 \tag{4.9}$$

On the other hand, as $1 \in I$ we have

$$b_1 = v_1 \leq \alpha v_{1^*} \tag{4.10}$$

Thus, we must have $v_2 \leq \frac{m}{m-1}\alpha v_{1^*} = \alpha_1 v_{1^*}$ or the second player would win the first slot.

Now let Γ be the set of players with value at least $\alpha_1 v_{1^*}$. Choose some constant γ. If $|\Gamma| < \gamma n$, then player 1^*'s lookahead utility for moving into slot one is at least $(1-\gamma)(1-\alpha_1)v_{1^*}c_1$. If player 1^* stays put, ignoring a repeat move for player 1^*, which occurs with probability $\frac{1}{n}$, player i^*'s utility in every other circumstance is at most

$$c_{1^*-1}v_{1^*} < \frac{1}{1-\alpha}c_{1^*}v_{1^*} < \frac{\beta}{1-\alpha}c_1 v_{1^*} \tag{4.11}$$

Since player 1^*'s utility is the same for both cases when a repeated move occurs and since we can choose β sufficiently small (i.e, $\beta < (1-\gamma)(1-\alpha)(1-\alpha_1)$), player 1^* will improve by moving into slot 1 in this case, contradicting the fact that we are at an equilibrium.

Thus, we may suppose $|\Gamma| > \gamma n$. Let $i_1 = \max_{i \in \Gamma} i$. Then the players in Γ contribute at least $\gamma n \alpha_1 v_{1^*} c_{i_1}$ to the social welfare. Take a constant δ and suppose that $c_{i_1} \geq \delta \frac{c_1}{n}$. Then the players in Γ would contribute at least $\gamma \delta \alpha_1 c_1 v_{1^*}$. Again, this is a constant fraction of social welfare that is contributed in the optimal allocation by player 1^* which, in turn, is a constant factor of the optimal social welfare of the slots in I. Thus, we would achieve a constant factor of the optimal social welfare.

So we may assume $c_{i_1} < \delta\frac{c_1}{n}$. Consider player i_1. His lookahead utility for staying in place, ignoring the case of a repeated move, is at most

$$c_{i_1-1}v_{i_1} \leq \frac{1}{1-\alpha}c_{i_1}v_{i_1} \leq \frac{1}{1-\alpha}\frac{\delta}{n}c_1v_{i_1} \leq \frac{1}{1-\alpha}\frac{\delta}{n}c_1v_{i^*} \tag{4.12}$$

We may assume that player $v_1 \leq (1-\epsilon)\alpha_1 v_{1^*}$, for some constant ϵ, otherwise we are done. Therefore, if player i_1 moves to slot 1 then he will earn at least $\epsilon c_1 v_{1^*}$ provided that player 1 makes the next move. This occurs with probability $1/n$, and so his total lookahead utility, ignoring a repeated move, is at least $\frac{\epsilon}{n}c_1v_{1^*}$. Thus by choosing $\delta \leq (1-\alpha)\epsilon$, it follows that the coordination ratio is constant in the average case model. □

CHAPTER 5
The Fisher Game

5.1 Introduction

As we saw in Chapter 1, general equilibrium is a fundamental concept in economics, tracing back to 1872 with the seminal work of Walras [175]. Traditionally, the focus of this theory has been upon *perfect competition*, where the number of buyers and sellers in the market are so huge that the contribution of any individual is infinitesimal. In particular, the participants are *price-takers*.

In practice, however, this assumption is unrealistic. This observation has motivated researchers to study markets where the players have an incentive to act strategically. A prominent example is the seminal work of Shapely and Shubik [154]. They defined *trading post games* for exchange markets and examined whether Nash equilibria there could implement competitive equilibrium prices and allocations. Another example, and a prime motivator of our research, is the Cournot-Walras market model introduced by Codognato and Gabszewicz [41] and Gabszewicz and Michel [70], which extends oligopolistic competition into the Walrasian setting. The importance of this model was demonstrated by Bonniseau and Florig [24] via a connection, in the limit, to traditional general equilibria models under the standard economic technique of agent *replication*. More recently, in the computer science community, Babaioff et al [12] extended Hurwicz's framework [85] to study the welfare of Walrasian markets acting through an auction mechanism.

Our interest is in how robust a pricing mechanism is against strategic manipulation. Specifically, our primary goal is to quantify the loss in social welfare due to price-making rather than price-taking behaviour. To do this, we define the *Price of Imperfect Competition (PoIC)* as the ratio of the social welfare at the worst Nash equilibrium to the social welfare at the perfectly-competitive Walrasian equilibrium.

Two remarks are pertinent here. First, we are interested in changes in the welfare produced by the market mechanism under the two settings of price-takers and price-makers. We are not interested in comparisons with the optimum social welfare, which requires the mechanism to possess the unrealistic power to perform total welfare redistribution. In particular, we are not concerned here with the *Price of Anarchy* or *Price of Stability*. Interestingly, though, the groundbreaking Price of Anarchy results of Johari and Tzitsiklis [91] on the proportional allocation mechanism for allocating one good (bandwidth) can be seen as the first Price of Imperfect Competition results. This is because in their setting the proportional allocation mechanism will produce optimal allocations in non-strategic settings; in contrast, for our markets, Walrasian equilibria can be arbitrarily poor in comparison to optimal allocations.

Second, in some markets the Price of Imperfect Competition may actually be larger than one. Thus, strategic manipulations by the agents can lead to improvements in social welfare! Indeed, we will exhibit examples where the social welfare increases by an arbitrarily large factor when the agents act strategically.

In this chapter, we analyze the Price of Imperfect Competition in Fisher markets with strategic buyers, a special case of the Cournot-Walras model. This scenario models the case of an oligopsonistic market, where the price-making power lies with the buyers rather than the sellers (as in an oligopoly).[1] Adsul et al. [3] study Fisher markets where buyers can lie about their preferences. They gave a complete characterization of its symmetric Nash equilibria (SNE) and showed that market equilibrium prices can be implemented at one of the SNE. Later Chen et. al. [37] studied *incentive ratios* in such markets to show that a buyer can gain no more than twice by strategizing in markets with linear, Leontief and Cobb-Douglas utility functions. In upcoming work, Branzei et al [29] study the Price of Anarchy in the game of Adsul et al. and prove polynomial lower and upper bounds for it. Furthermore, they show Nash equilibria always exist.

In the above games (and the Fisher model itself), only the sellers have an intrinsic utility for money. In contrast, we postulate that buyers (and not just sellers) have utility for money. Thus, buyers may also benefit by saving money for later use. This incentivizes buyers to withhold money from the market. This defines our *Fisher Market Game*, where agents strategize on the amount of money they wish to spend, and obtain utility one from each unit of saved money. Contrary to the bound of two on gains when strategizing on utility functions [37], we observe that strategizing on money may facilitate unbounded gains (see Section 5.7.1). These

[1] The importance of oligopsonies was recently highlighted by the price-fixing behaviour of massive technology companies in San Francisco.

incentives can induce large variations between the allocations produced at a Market equilibrium and at a Nash equilibrium. Despite this, we prove the Price of Imperfect Competition is at least $\frac{1}{2}$ for Fisher markets when the buyers utility functions belong to the utililty class of Constant Elasticity of Substitution (CES) with the weak gross substitutability property – this class includes linear and Cobb-Douglas functions.

5.1.1 Overview of Chapter

In Section 5.2, we define the Fisher Game and present our welfare metrics. In Section 5.3, we prove that Price of Imperfect Competition is at least $\frac{1}{2}$, for CES utilities which satisfy the weak gross substitutability property. In Section 5.4, we apply the economic technique of replication to demonstrate that, for linear utilities, the PoIC bound improves as the level of competition in the market increases. In Section 5.5, we turn our attention to the question of existence of Nash equilibria. We establish that Nash equilibria exist for the subclass of Cobb-Douglas utilities. However, they need not exist for all CES utilities. In particular, Nash equilibria need not exist for linear utilities. To address this possibility of non-existence, in Section 5.6, we examine the dynamics of the linear Fisher Game and provide logarithmic welfare guarantees. In Section 5.7, we provide examples of Fisher Games to show that the PoIC can range from below 1 to unboundedly high. In Section 5.8, we demonstrate that for another game, the Proportional Allocation Mechanism, the PoIC can also be arbitrarily small.

5.2 Preliminaries

We now define the Fisher Game where agents strategize on how much money to spend.

5.2.1 The Fisher Game.

Recall the Fisher Market Model from Section 2.2.1. An implicit assumption within this model is that money has an intrinsic value *to the sellers*, stemming from its potential use outside of the market or at a later date. Thus, money is not just a numéraire. In our model, we assume this intrinsic value applies to all market participants including the buyers. This assumption induces a strategic game in which the buyers may have an incentive to save some of their money.

This *Fisher Game* is a special case of the general Cournot-Walras game introduced by Codognato, Gabszewicz, and Michel ([41], [70]). Here the buyers can choose some strategic amount of money $s_i < m_i$ to bring to the market, which will affect their budget constraint. They gain utility both from the resulting market equilibria (with s_i substituted for m_i) and from the money they withhold from the market. Observe, in the Fisher market model, the sellers have no value for the goods in the market. Thus, in the corresponding game, they will place all their goods on sale as their only interest is in money. (Equivalently, we may assume the sellers are non-strategic.)

Thus, we are in an oligopsonistic situation where buyers have indirect price-making power. The set of strategies available to buyer i is $M_i = \{s \geq 0 \mid s \leq m_i\}$. When each buyer decides to spend $s_i \in M_i$, then $\boldsymbol{p}(\boldsymbol{s})$ and $\boldsymbol{x}(\boldsymbol{s})$ are the prices and allocations, respectively, produced by the Fisher market mechanism. These can be determined from the Eisenberg-Gale program (2.13) by substituting s_i for m_i. Thus, the total payoff to buyer i is

$$T_i(\boldsymbol{s}) = U_i(\boldsymbol{x}_i(\boldsymbol{s})) + (m_i - s_i) \tag{5.1}$$

Our primary tool to analyze the Fisher Game is via the standard solution concept of a Nash equilibrium. A strategy profile $\boldsymbol{s}$ is said to be a *Nash equilibrium* if no player gains by deviating unilaterally. Formally, $\forall i \in \mathcal{B},\ T_i(\boldsymbol{s}) \geq T_i(s', \boldsymbol{s}_{-i}),\ \forall s' \in M_i$. For the market game defined on market $\mathcal{M}$, let $NE(\mathcal{M})$ denote its set of NE strategy profiles.

The incentives in the Fisher Game can be high. In particular, in Section 5.7.1,we show that for any $L \geq 0$, there is a market with linear utility functions where an agent improve his payoff by a multiplicative factor of L by acting strategically.

5.2.2 The Price of Imperfect Competition.

The social welfare of a strategy is the aggregate payoff of both buyers and sellers. At a state $\boldsymbol{s}$, with prices $\boldsymbol{p} = \boldsymbol{p}(\boldsymbol{s})$ and allocations $\boldsymbol{x} = \boldsymbol{x}(\boldsymbol{s})$, the social welfare is:

$$\mathcal{W}(\boldsymbol{s}) \;=\; \sum_{i\in\mathcal{B}}(U_i(\mathbf{x_i}) + m_i - s_i) + \sum_{j\in\mathcal{G}} p_j \;=\; \sum_{i\in\mathcal{B}} U_i(\mathbf{x_i}) + \sum_{i\in\mathcal{B}} m_i \tag{5.2}$$

Note, here, that the cumulative payoff of sellers is $\sum_{j\in\mathcal{G}} p_j = \sum_{i\in\mathcal{B}} s_i$.

The focus of this chapter is how strategic manipulations of the market mechanism affect the overall social welfare. Thus, we must compare the social welfare of the strategic Nash equilibrium to that of the unstrategic market equilibrium where all buyers simply put all of their money onto the market. This latter equilibrium is the *Walrasian equilibrium (WE)*. This comparison gives rise to a welfare ratio, which we term the *Price of Imperfect Competition (PoIC)*, the ratio of the minimum welfare amongst strategic Nash equilibria in the market game to the welfare of the unstrategic Walrasian equilibrium. Formally, for a given market $\mathcal{M}$,

$$\text{PoIC}(\mathcal{M}) = \min_{\boldsymbol{s} \in NE(\mathcal{M})} \frac{\mathcal{W}(\boldsymbol{s})}{\mathcal{W}(\boldsymbol{m})}$$

Thus the Price of Imperfect Competition is a measure of how robust, with respect to social welfare, the market mechanism is against oligopsonist behaviour. Observe that the Price of Imperfect Competition could be either greater or less than 1. Indeed, the example in Section 5.7.1 shows that a Nash Equilibrium may produce arbitrarily higher welfare than a Walrasian Equilibrium. Of course, one may expect that welfare falls when the mechanism is gamed and, in Section 5.7.2, we do present an example where the welfare at a Nash Equilibrium is slightly lower than at the Walrasian Equilibrium. This leads to the question of whether the welfare at a Nash can be much worse than at a market equilibrium. We will show that the answer is no; a Nash always produces at least a constant factor of the welfare of a market equilibrium.

5.3 Bounds on the Price of Imperfect Competition

In this section we establish bounds on the PoIC for the Fisher Game for CES utilities with $0 < \rho \leq 1$ and for Cobb-Douglas utilities. The example in Section 5.7.1 shows that there is no upper bound on PoIC for the Fisher Game. Thus, counter-intuitively, even for linear utilities, it may be extremely beneficial to society if the players are strategic.

In the rest of this section, we demonstrate a lower bound of $\frac{1}{2}$ on the PoIC. This result distinguishes the Fisher Game from other strategic market models. For example, consider the case of the Proportional Allocation Mechanism applied over a multi-good market (see Feldman et al. [63] for details on this application). In Section 5.8, we show that the PoIC may then approach zero in the proportional allocation mechanism with savings. Thus the Fisher Game is, in a sense, more resilient to strategic play than other mechanisms.

So take a market with Cobb-Douglas or CES utility functions (where $0 < \rho \leq 1$). We assume that the market is non-trivial in the sense that for each buyer i, there is some good j such that $u_{ij} > 0$ and also fore each good j, there is some buyer i for which $u_{ij} > 0$. The key to proving the factor $\frac{1}{2}$ lower bound on the PoIC is the following lemma showing the monotonicity of prices.

Lemma 3. *Given two strategic allocations of money* $\mathbf{s}^* \leq \mathbf{s}$, *then the corresponding equilibrium prices satisfy* $\boldsymbol{p}^* \leq \boldsymbol{p}$, *where* $\boldsymbol{p}^* = \boldsymbol{p}(\boldsymbol{s}^*)$ *and* $\boldsymbol{p} = \boldsymbol{p}(\boldsymbol{s})$.

Proof. We break the proof up into three classes of utility function.

`(i) Cobb-Douglas Utilities`

The case of Cobb-Douglas utility functions is simple. To see this, recall a result of Eaves [49]. He showed that, when buyer i spends s_i, the prices and allocations for the Fisher market are given by

$$p_j = \sum_i u_{ij} s_i \qquad x_{ij} = \frac{u_{ij} s_i}{\sum_k u_{kj} s_k} \tag{5.3}$$

It follows that if strategic allocations of money increase, then so must prices.

`(ii) CES Utilities with` $0 < \rho < 1$

Recall that market equilibria for CES Utilities can be calculated via the Eisenberg-Gale convex program (2.13). From the KKT conditions of this program, where p_j is the dual variable of the budget constraint, we observe that:

$$\begin{aligned} &\forall j, && p_j > 0 \Rightarrow \textstyle\sum_i x_{ij} = 1 \\ &\forall (i,j), && \tfrac{s_i u_{ij}}{U_i(\boldsymbol{x}_i)^\rho x_{ij}^{1-\rho}} \le p_j \quad \text{and} \quad x_{ij} > 0 \Rightarrow \tfrac{s_i u_{ij}}{U_i(\boldsymbol{x}_i)^\rho x_{ij}^{1-\rho}} = p_j \end{aligned} \tag{5.4}$$

Claim 2. *If players have CES utilities with* $0 < \rho < 1$ *and* $\sigma_i > 0$*, then* $x_{ij} > 0, \ \forall (i,j)$ *with* $u_{ij} > 0$.

Proof. Consider the derivative of U_i with respect to x_{ij} as $x_{ij} \to 0$:

$$\lim_{x_{ij} \to 0} \frac{\partial U_i(\boldsymbol{x}_i)}{\partial x_{ij}} \quad = \quad \lim_{x_{ij} \to 0} \frac{u_{ij} U_i(\boldsymbol{x}_i)^{1-\rho}}{x_{ij}^{1-\rho}} \quad = \quad +\infty \tag{5.5}$$

The claim follows since $p_j \le \sum_i s_i$ and is, thus, finite. □

We may now proceed by contradiction. Suppose $\exists k$ s.t. $p_k < p_k^*$. Choose a good j such that $\frac{p_j}{p_j^*}$ is minimal and therefore less than 1, by assumption. Take any player i such that $u_{ij} > 0$. By the above claim, we have $x_{ij}, x_{ij}^* > 0$. Consequently, by the KKT conditions (5.4), we have:

$$\frac{u_{ij}}{p_j x_{ij}^{1-\rho}} = \frac{U_i(\boldsymbol{x}_i)^\rho}{s_i} \qquad \text{and} \qquad \frac{u_{ij}}{p_j^* x_{ij}^{*1-\rho}} = \frac{U_i(\boldsymbol{x}_i^*)^\rho}{s_i^*} \tag{5.6}$$

Taking a ratio gives:

$$\frac{p_j x_{ij}^{1-\rho}}{p_j^* x_{ij}^{*1-\rho}} = \frac{U_i(\boldsymbol{x}_i^*)^\rho s_i}{U_i(\boldsymbol{x}_i)^\rho s_i^*} \tag{5.7}$$

Indeed, this equation also holds for every good $t \in \mathcal{G}$ with $u_{it} > 0$. Next consider the following two cases:

Case 1: $x_{ij} \leq x_{ij}^*$ for some player i.

From (5.7) we must then have that $U_i(\boldsymbol{x}_i) > U_i(\boldsymbol{x}_i^*)$. However, by the minimality of $\frac{p_j}{p_j^*}$, and since (5.7) holds for every $t \in \mathcal{G}$ with $u_{it} > 0$, we obtain $x_{it} \leq x_{it}^*$ for all such t. This implies $U_i(\boldsymbol{x}_i) \leq U_i(\boldsymbol{x}_i^*)$, a contradiction.

Case 2: $x_{ij} > x_{ij}^*$ for every player i.

Since $p_j^* > p_j$, we must have $p_j^* > 0$. By (5.4) it follows that $\sum_i x_{ij}^* = 1$. But now we obtain the contradiction that demand must exceed supply as $\sum_i x_{ij} > \sum_i x_{ij}^* = 1$.

`(iii) Linear Utilities`

We begin with some notation. Let $S_i = \{j \in \mathcal{G} : x_{ij} > 0\}$ be the set of goods purchased by buyer i at strategy $\boldsymbol{s}$. Let $\beta_{ij} = \frac{u_{ij}}{p_j}$ be the *rate-of-return* of good j for buyer i at prices $\mathbf{p}$. Let $\beta_i = \max_{j \in \mathcal{G}} \beta_{ij}$ be the *bang-for-buck* buyer i can obtain at prices $\mathbf{p}$. It can be seen from the KKT conditions of the Eisenberg-Gale program (2.13) that at $\{\boldsymbol{p}, \boldsymbol{x}\}$, every good $j \in S_i$ will have a rate-of-return equal to the bang-for-buck (see, for example, [172]). Similarly, let S_i^*, β_i^* be correspondingly defined for strategy s^*.

Note that, assuming for each good j, $\exists i,\ u_{ij} > 0$, we have that $\boldsymbol{p}, \boldsymbol{p}^* > 0$. Thus, we can partition the goods into groups based on the *price ratios* $\frac{p_j^*}{p_j}$. Suppose there are k distinct price ratios over all the goods (thus $k \leq g$), then partition the goods into k groups, say $\mathcal{G}_1, \ldots, \mathcal{G}_k$ such that all the goods in a group have the same ratio. Let the ratio in group j be λ_j and let $\lambda_1 < \lambda_2 < \cdots < \lambda_k$. Thus $\mathcal{G}_1$ are the goods whose prices have fallen the most (risen the least) and $\mathcal{G}_k$ are the goods whose prices have fallen the least (risen the most).

Let $\mathcal{I}_k = \{i : \exists j \in \mathcal{G}_k, x_{ij} > 0\}$ and $\mathcal{I}_k^* = \{i : \exists j \in \mathcal{G}_k, x_{ij}^* > 0\}$. Thus $\mathcal{I}_k$ and $\mathcal{I}_k^*$ are the collections of buyers that purchase goods in $\mathcal{G}_k$ in each of the allocations. Take any buyer $i \in \mathcal{I}_k^*$; so there is some good $j \in S_i^* \cap \mathcal{G}_k$.

If $S_i \cap \bigcup_{\ell=1}^{k-1} \mathcal{G}_\ell \neq \emptyset$ then buyer i would not desire good j at prices p_j^*. To see this, take a good $j' \in S_i \cap \bigcup_{\ell=1}^{k-1} \mathcal{G}_\ell$. Then $\beta_{ij'} = \beta_i \geq \beta_{ij}$. Therefore

$$\begin{aligned} \beta_i^* &\geq \frac{u_{ij'}}{p_{j'}^*} \geq \frac{u_{ij'}}{\lambda_{k-1} \cdot p_{j'}} > \frac{u_{ij'}}{\lambda_k \cdot p_{j'}} \\ &= \frac{1}{\lambda_k} \cdot \frac{u_{ij'}}{p_{j'}} \geq \frac{1}{\lambda_k} \cdot \frac{u_{ij}}{p_j} \\ &= \frac{u_{ij}}{p_j^*} = \beta_i^* \end{aligned} \tag{5.8}$$

This contradiction tells us that $S_i \subseteq \mathcal{G}_k$ and $\mathcal{I}_k^* \subseteq \mathcal{I}_k$. It follows that $\cup_{i \in \mathcal{I}_k^*} S_i \subseteq \mathcal{G}_k$. Putting this together, we obtain that

$$\sum_{i \in \mathcal{I}_k^*} s_i \leq \sum_{i \in \mathcal{I}_k} s_i \leq \sum_{j \in \mathcal{G}_k} p_j \tag{5.9}$$

Now recall that all goods must be sold by the market mechanism (as $\boldsymbol{p}, \boldsymbol{p}^* > 0$). Thus the buyers $\mathcal{I}_k^*$ must be able to afford all of the goods in $\mathcal{G}_k$. Thus

$$\sum_{i \in \mathcal{I}_k^*} s_i^* \geq \sum_{j \in \mathcal{G}_k} p_j^* = \lambda_k \cdot \sum_{j \in \mathcal{G}_k} p_j \tag{5.10}$$

But $s_i^* \le s_i$ for all i. Consequently, Inequalities (5.9) and (5.10) imply that $\lambda_k \le 1$. Thus no price in $\boldsymbol{p}^*$ can be higher than in $\boldsymbol{p}$. $\square$

First we use Lemma 3 to provide lower bounds on the individual payoffs.

Lemma 4. *Let s_i be a best response for agent i against the strategies $\boldsymbol{s}_{-i}$. Then $T_i(\boldsymbol{s}) \ge \max(\hat{U}_i, m_i)$, where $\hat{U}_i$ is her utility at the Walrasian equilibrium.*

Proof. Clearly $T_i(\boldsymbol{s}) \ge m_i$, otherwise player i could save all her money and achieve a payoff of m_i. For $T_i(\boldsymbol{s}) \ge \hat{U}_i$, let $\boldsymbol{p} = \boldsymbol{p}(\boldsymbol{m})$ and $\boldsymbol{x} = \boldsymbol{x}(\boldsymbol{m})$ be the prices and allocation at Walrasian equilibrium. If buyer i decides to spend all his money when the others play $\boldsymbol{s}_{-i}$, the resulting equilibrium prices will be less than $\boldsymbol{p}$, by Lemma 3. Therefore, she can afford to buy bundle $\boldsymbol{x}_i$. Thus, her best response payoff must be at least $\hat{U}_i$. $\square$

It is now easy to show the lower bound on the Price of Imperfect Competition.

Theorem 14. *In the Fisher Game, with Cobb-Douglas or CES utilities ($0 < \rho \le 1$), we have PoIC $\ge \frac{1}{2}$. That is, $\mathcal{W}(\boldsymbol{s}^*) \ge \frac{1}{2}\mathcal{W}(\boldsymbol{m})$, for any Nash equilibrium $\boldsymbol{s}^*$.*

Proof. Let $\boldsymbol{p}^* = \boldsymbol{p}(\boldsymbol{s}^*)$ and $\boldsymbol{x}^* = \boldsymbol{x}(\boldsymbol{s}^*)$. Let $\boldsymbol{p}$ and $\boldsymbol{x}$ be the Walrasian equilibrium prices and allocations, respectively. At the Nash equilibrium $\boldsymbol{s}^*$ we have $T_i(\boldsymbol{s}^*) \ge \max(m_i, U_i(\boldsymbol{x}_i))$ for each player i, by Lemma 4. Thus, we obtain:

$$2\sum_i T_i(\boldsymbol{s}^*) \ge \sum_i U_i(\boldsymbol{x}_i) + \sum_i m_i \tag{5.11}$$

Therefore $\mathcal{W}(\boldsymbol{s}^*) \ge \frac{1}{2}\mathcal{W}(\boldsymbol{m})$, as desired. $\square$

5.4 Social Welfare and the Degree of Competition

In this section, we examine how the welfare guarantee improves with the degree of competition in the market. To model the degree of competition, we apply a common technique in the economics literature, namely *replication* [154]. In a replica economy, we take each buyer type in the market and make N duplicates (the budgets of each duplicate is a factor N smaller than that of the original buyer). The *degree of competition* in the resultant market is N. We now consider the Fisher Game with linear utility functions and show how the lower bound on Price of Imperfect Competition improves with N.

Theorem 15. *Let $\boldsymbol{s}^*$ be a NE in a market with degree of competition N. Then*

$$\mathcal{W}(\boldsymbol{s}^*) \geq (1 - \frac{1}{N+1}) \cdot \mathcal{W}(\boldsymbol{m}) \tag{5.12}$$

In order to prove Theorem 15, we need a better understanding of how prices adjust to changes in strategy under different degrees of competition. Towards this goal, we need the following two lemmas.

Lemma 5. *Given an arbitrary strategic money allocation $\boldsymbol{s}$. If player i increases (resp. decreases) her spending from s_i to $(1+\delta)s_i$ then the price of any good increases (resp. decreases) by at most a factor of $(1+\delta)$.*

Proof. We focus on the case of increase; the argument for the decrease case is analogous. Suppose all players increase their strategic allocation by a factor of $(1+\delta)$. Then the allocations to all players would remain the same by the market mechanism and all prices would be scaled up by a factor of $(1+\delta)$. Then suppose each player

$k \neq i$ subsequently lowers its money allocation back down to the original amount s_k. By Lemma 3, no price can now increase. The result follows. □

Lemma 6. *Given an arbitrary strategic money allocation* $\boldsymbol{s}$ *in a market with degree of competition* N*. Let buyer* i *be the duplicate player of her type with the smallest money allocation* s_i*. If she increases her spending to* $(1 + N \cdot \delta)s_i$ *then the price of any good increases by at most a factor* $(1 + \delta)$.

Proof. We utilize the symmetry between the N identical players. Let players $i_1 = i, i_2, ..., i_N$ be the replicas identical to player i. If each of these players increased their spending by a factor of $(1 + \delta)$ then, by Lemma 5, prices would go up by at most a factor $(1 + \delta)$. From the market mechanism's perspective, this is equivalent to player i increasing her strategic allocation to $s_i + \delta \cdot \sum_k s_{i_k}$. But this is greater than $(1 + N \cdot \delta)s_i$. Thus, by Lemma 3, the new prices are larger by a factor of at most $(1 + \delta)$. □

Now let $\boldsymbol{x} = \boldsymbol{x}(\boldsymbol{m})$ and $\boldsymbol{x}^* = \boldsymbol{x}(\boldsymbol{s}^*)$. Since we have rational inputs, $\mathbf{x}$ and $\mathbf{x}^*$ must be rational [89]. Therefore, by appropriately duplicating the goods and scaling the utility coefficients, we may assume that there is exactly one unit of each good and that both $\mathbf{x}$ and $\mathbf{x}^*$ are $\{0, 1\}$-allocations. Recall from the proof of Lemma 3 our definition of S_i, S_i^* and β_i, β_i^*. Under this assumption, $S_i = \{j \in \mathcal{G} : x_{ij} = 1\}$ and similarly for S_i^*. We are now ready to prove the following welfare lemma.

Lemma 7. *For any Nash equilibrium* $\{\mathbf{s}^*, \mathbf{p}^*, \mathbf{x}^*\}$ *and any Walrasian equilibrium* $\{\mathbf{s} = \mathbf{m}, \mathbf{p}, \mathbf{x}\}$*, we have*

$$\sum_{i \in \mathcal{B}} \sum_{j \in S_i^*} u_{ij} \geq \left(1 - \frac{1}{N}\right) \cdot \sum_{i \in \mathcal{B}} \sum_{j \in S_i} u_{ij} \tag{5.13}$$

Proof. To prove the lemma we show that total utility produced by goods at NE, after scaling by a factor $\frac{N}{N-1}$, is at least as much as the utility they produce at the Walrasian equilibrium. We do this by partitioning goods into the sets S_i. We then notice that for each good, the player who receives it at NE must receive utility from it in excess of the price he paid for it. In many cases, this price is more than the utility of the player who receives it in Walrasian equilibrium and we are done. Otherwise we will set up a transfer system where players in NE who receive more utility for the good than the price paid for it transfer some of this excess utility to players who need it. This will ultimately allow us to reach the desired inequality.

For the rest of this proof, without loss of generality, we will restrict our attention to Nash equilibria where each identical copy of a certain type of player has the same strategy. We are able to do this as the market could treat the sum of these copies as a single player and thus we are able to manipulate the allocations between these players without changing market prices or the total utility derived from market allocations. Thus if our argument holds for Nash equilibria where identical players have the same strategy, it will also hold for heterogeneous Nash equilibria. Now take any player i. There are two cases:

Case 1: $s_i^* = m_i$.

By Lemma 3, we know that

$$\sum_{j \in S_i^* \cap S_i} p_j^* \leq \sum_{j \in S_i^* \cap S_i} p_j \tag{5.14}$$

Therefore, by the assumption that $s_i^* = m_i$, we have

$$\sum_{j \in S_i \setminus S_i^*} p_j = m_i - \sum_{j \in S_i^* \cap S_i} p_j = s_i^* - \sum_{j \in S_i^* \cap S_i} p_j \leq s_i^* - \sum_{j \in S_i^* \cap S_i} p_j^* = \sum_{j \in S_i^* \setminus S_i} p_j^* \tag{5.15}$$

Thus buyer i spends more on $S_i^* \setminus S_i$ than she did on $S_i \setminus S_i^*$. But, by Lemma 3, she also receives a better bang-for-buck on $S_i^* \setminus S_i$ than on $S_i \setminus S_i^*$, as $\beta_i^* \geq \beta_i$ (Lemma 3). Let $\beta_i^* = 1 + \epsilon_i^*$. Thus, at the Nash equilibrium, her total utility on $S_i^* \setminus S_i$ is

$$\sum_{j \in S_i^* \setminus S_i} u_{ij} \;=\; \sum_{j \in S_i^* \setminus S_i} \beta_i^* \cdot p_j^* \;=\; (1 + \epsilon_i^*) \cdot \sum_{j \in S_i^* \setminus S_i} p_j^* \tag{5.16}$$

Of this utility, buyer i will allocate p_j^* units of utility to each item $j \in S_i^* \setminus S_i$. The remaining $\epsilon_i^* \cdot p_j^*$ units of utility derived from good j is reallocated to goods in $S_i \setminus S_i^*$.

Consider the goods in S_i. Clearly goods in $S_i \cap S_i^*$ contribute the same utility to both the Walrasian equilibrium and the Nash equilibrium. So take the items in $S_i \setminus S_i^*$. The buyers of these items at NE have obtained at least $\sum_{j \in S_i \setminus S_i^*} p_j^*$ units of utility from them (as $\beta_d^* \geq 1$, $\forall d$). In addition, buyer i has reallocated $\epsilon_i^* \cdot \sum_{j \in S_i^* \setminus S_i} p_j^*$ to goods in $S_i \setminus S_i^*$. So the total utility allocated to goods in $S_i \setminus S_i^*$ is

$$\begin{aligned}
\sum_{j \in S_i \setminus S_i^*} p_j^* + \epsilon_i^* \cdot \sum_{j \in S_i^* \setminus S_i} p_j^* \;&\geq\; \sum_{j \in S_i \setminus S_i^*} p_j^* + \epsilon_i^* \cdot \sum_{j \in S_i \setminus S_i^*} p_j^* \;=\; (1 + \epsilon_i^*) \cdot \sum_{j \in S_i \setminus S_i^*} p_j^* \\
&=\; \beta_i^* \cdot \sum_{j \in S_i \setminus S_i^*} p_j^* \;\geq\; \sum_{j \in S_i \setminus S_i^*} u_{ij}
\end{aligned} \tag{5.17}$$

Here the first inequality follows by (5.15) and the final inequality follows as $\beta_i^* \geq \frac{u_{ij}}{p_j^*}$, for any good $j \notin S_i^*$. Thus the reallocated utility on S_i at NE is greater than the utility it provides in the Walrasian equilibrium (even without scaling by $\frac{N}{N-1}$).

Case 2: $s_i^* < m_i$.

Suppose buyer i increases her spending from s_i^* to $(1 + N \cdot \delta) \cdot s_i^*$. Then the prices of the goods she buys increase by at most a factor $(1 + \delta)$ by Lemma 6. Thus her utility changes by

$$(m_i - (1 + \delta \cdot N) \cdot s_i^*) + s_i^* \cdot \beta_i^* \cdot \frac{1 + N \cdot \delta}{1 + \delta} - (m_i - s_i^*) - s_i^* \cdot \beta_i^* \ \leq \ 0 \qquad (5.18)$$

where the inequality follows as s^* is a Nash equilibrium. This simplifies to

$$s_i^* \cdot \left(-\delta \cdot N + \beta_i^* \cdot (\frac{1 + N \cdot \delta}{1 + \delta} - 1) \right) \leq 0 \qquad (5.19)$$

Now suppose (i) $s_i^* = 0$. In this case we must have $u_{ij}/p_j^* \leq 1$ for every good j. To see this, we argue by contradiction. Suppose $u_{ij}/p_j^* = 1 + \epsilon$ for some good j. Notice that if player i changes s_i^* to γ the price of good j can go up by at most γ as we know each price increases by Lemma 3 and the sum of all prices is at most γ higher (by the market conditions). Thus, if player i puts $\gamma < \epsilon$ money onto the market then good j will still have bang-for-buck greater than 1 and so player i will gain more utility than the loss of savings. Thus, s_i^* cannot be an equilibrium, a contradiction.

Thus $u_{ij} \leq p^*_j \leq u_{i^*j}$ where i^* is the player who receives good j at NE. Therefore this player obtains more utility from good j than player i did in the Walrasian equilibrium, even without scaling or a utility transfer.

On the other hand, suppose (ii) $s^*_i > 0$. This can only occur if we have both $\beta^*_i \geq 1$ and

$$\beta^*_i \cdot \frac{(N-1)\cdot\delta}{1+\delta} \leq \delta \cdot N \tag{5.20}$$

Therefore $1 \leq \beta^*_i \leq (1+\delta)\cdot(1+\frac{1}{N-1})$. Since this holds for all δ, as we take $\delta \to 0$ we must have $\beta^*_i \leq \frac{N}{N-1}$. Thus $\frac{u_{ij}}{p^*_j} \leq \frac{N}{N-1}$ for every good j. Thus if we multiply the utility of the player receiving good j in the Nash equilibrium by $\frac{N}{N-1}$ he will be getting more utility from it than player i did in the Walrasian equilibrium. □

Proof of Theorem 15. Given the other buyers strategies $\boldsymbol{s}^*_{-\mathbf{i}}$ suppose buyer i sets $s_i = m_i$. Then, by Lemma 3, prices cannot be higher for $(m_i, \boldsymbol{s}^*_{-\mathbf{i}})$ than at the Walrasian equilibrium $\boldsymbol{p}(\boldsymbol{m})$. Therefore, by selecting $s_i = m_i$, buyer i could afford to buy the entire bundle S_i at the resultant prices. Consequently, her best response strategy $\boldsymbol{s}^*_i$ must offer at least that much utility. This is true for each buyer, so we have

$$\sum_{i\in\mathcal{B}}\left((m_i - s^*_i) + \sum_{j\in\mathcal{G}} u_{ij}\cdot x^*_{ij}\right) \geq \sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{G}} u_{ij}\cdot x_{ij} \tag{5.21}$$

Thus

$$\begin{aligned}\mathcal{W}(\boldsymbol{s}^*) &= \sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{G}} u_{ij}\cdot x^*_{ij} + \sum_{i\in\mathcal{B}} m_i \;=\; \sum_{i\in\mathcal{B}}\left((m_i - s^*_i) + \sum_{j\in\mathcal{G}} u_{ij}\cdot x^*_{ij}\right) + \sum_{i\in\mathcal{B}} s^*_i \\ &\geq \sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{G}} u_{ij}\cdot x_{ij} + \sum_{i\in\mathcal{B}} s^*_i \end{aligned} \tag{5.22}$$

On the other hand, Lemma 7 implies that

$$\mathcal{W}(\boldsymbol{s}^*) \;=\; \sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{G}} u_{ij}\cdot x^*_{ij} + \sum_{i\in\mathcal{B}} m_i \;\geq\; \left(1-\frac{1}{N}\right)\cdot \sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{G}} u_{ij}\cdot x_{ij} + \sum_{i\in\mathcal{B}} m_i$$

Taking a convex combination of Inequalities (5.22) and (5.23) gives

$$\begin{aligned}
\mathcal{W}(\boldsymbol{s}^*) \;&\geq\; \left(\alpha\cdot(1-\frac{1}{N}) + (1-\alpha)\right)\cdot \sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{G}} u_{ij}\cdot x_{ij} + \alpha\cdot\sum_{i\in\mathcal{B}} m_i + (1-\alpha)\cdot\sum_{i\in\mathcal{B}} s^*_i \\
&\geq\; \left(\alpha\cdot(1-\frac{1}{N}) + (1-\alpha)\right)\cdot \sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{G}} u_{ij}\cdot x_{ij} + \alpha\cdot\sum_{i\in\mathcal{B}} m_i \\
&=\; \left(1-\frac{\alpha}{N}\right)\cdot \sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{G}} u_{ij}\cdot x_{ij} + \alpha\cdot\sum_{i\in\mathcal{B}} m_i
\end{aligned} \tag{5.23}$$

Thus plugging $\alpha = \frac{N}{N+1}$ in (5.23) gives

$$\mathcal{W}(\boldsymbol{s}^*) \;\geq\; \left(1-\frac{1}{N+1}\right)\cdot\left(\sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{G}} u_{ij}\cdot x_{ij} + \sum_{i\in\mathcal{B}} m_i\right) \;=\; \left(1-\frac{1}{N+1}\right)\cdot\mathcal{W}(\boldsymbol{m}) \tag{5.24}$$

This completes the proof. □

5.5 Existence of Nash Equilibria

We have demonstrated bounds for the Price of Imperfect Competition in the Fisher Game under both CES and Cobb-Douglas utilities. However, these welfare results only apply to strategies that are Nash equilibria. In this section, we prove that Nash equilibria exist for the Cobb-Douglas case, but need not exist for linear utilities. For games without Nash equilibria, we may still recover some welfare guarantees; we show this in Section 5.6, by examining the dynamics of the Fisher Game with linear utilities.

5.5.1 Cobb-Douglas Utility Functions

We will prove that a Nash equilibrium always exists for Fisher Games with Cobb-Douglas utilities as long as each good provides utility for at least two players.[2]

Recall that $T_i(\boldsymbol{s})$ is player i's total utility at strategy profile $\boldsymbol{s}$. The first step in this proof is to show that T_i is a concave function with respect to s_i when $\boldsymbol{s}_{-i}$ is fixed.

Lemma 8. *T_i is a concave function of s_i.*

Proof. First, it is enough for us to consider the component of the utility from the market, U_i (as the utility from saving money is always concave). Recall that from (5.3), we have $y_{ij} = x_{ij} \cdot p_j = s_i \cdot u_{ij}$. Thus, we can easily express U_i as a function of s_i as:

$$U_i \;=\; \prod_j x_{ij}^{u_{ij}} \;=\; \prod_j \left(\frac{s_i \cdot u_{ij}}{\tilde{p}_j + s_i u_{ij}} \right)^{u_{ij}} \tag{5.25}$$

Here $\tilde{p}_j = \sum_{k \neq i} y_{kj}$. We get the second equality simply by writing each x_{ij} as $\frac{y_{ij}}{p_j}$. Now, note that $\prod_j u_{ij}^{u_{ij}}$ is just a positive constant and so does not affect concavity. Also, $\prod_j s_i^{u_{ij}} = s_i$ by our assumption that $\sum_j u_{ij} = 1$. Thus it is enough to show that the following is concave:

$$\tilde{U}_i = \frac{s_i}{\prod_j (\tilde{p}_j + s_i u_{ij})^{u_{ij}}}. \tag{5.26}$$

[2] In the absence of this assumption, it is possible for a player who is a monopsonist of a single good to continually decrease their strategic allocation, trivially precluding the possibility of an equilibrium.

Taking derivatives give us:

$$\tilde{U}_i' = \frac{\prod_j (\tilde{p}_j + s_i u_{ij})^{u_{ij}} - s_i \sum_k u_{ik}^2 (\tilde{p}_k + s_i u_{ik})^{(u_{ik}-1)} \prod_{j \neq k} (\tilde{p}_j + s_i u_{ij})^{u_{ij}}}{\prod_j (\tilde{p}_j + s_i u_{ij})^{2u_{ij}}} \tag{5.27}$$

Notice that the numerator simplifies considerably, if we take advantage of the fact that $\sum_j u_{ij} = 1$ to rewrite it as:

$$\begin{aligned} & \sum_k u_{ik} \prod_j (\tilde{p}_j + s_i u_{ij})^{u_{ij}} - s_i \sum_k u_{ik}^2 (\tilde{p}_k + s_i u_{ik})^{(u_{ik}-1)} \prod_{j \neq k} (\tilde{p}_j + s_i u_{ij})^{u_{ij}} \\ = \; & \sum_k \tilde{p}_k (\tilde{p}_k + s_i u_{ik})^{(u_{ik}-1)} \prod_{j \neq k} (\tilde{p}_j + s_i u_{ij})^{u_{ij}} \end{aligned} \tag{5.28}$$

Thus, we can simplify to

$$\tilde{U}_i' = \sum_k \frac{\tilde{p}_k}{(\tilde{p}_k + s_i u_{ik}) \prod_j (\tilde{p}_j + s_i u_{ij})^{u_{ij}}} \tag{5.29}$$

But this is clearly a decreasing function of s_i and so $\tilde{U}_i$ is concave. □

We are now ready to prove the existence of an equilibrium.

Theorem 16. *If for every good at least two players have positive utility for that good, then a Nash equilibrium of the strategic game exists.*

Proof. This proof is similar in structure to that of [63]. Let $\Gamma = (\boldsymbol{U}, \boldsymbol{m})$ be the original market game. For each $\epsilon > 0$, we define the epsilon-market as Γ_ϵ. This market has all of the original players and goods, but will limit the strategy sets of each player by forcing them to put at least ϵ of their money on the market.

It is easy to see that in the epsilon version of the game, utilities are continuous with respect to the strategic variable. This follows from (5.25). Also, by Lemma 8, we see that the function T_i with respect to s_i is concave. Applying Rosen's theorem [137] we get that a market equilibrium must exist for each epsilon market. Let $\boldsymbol{s}^*_\epsilon$ be this equilibrium.

Notice that, since the strategy sets are compact, there must be a limit point to $\boldsymbol{s}^*_\epsilon$ as $\epsilon \to 0$. Call this point $\boldsymbol{s}^*$. Clearly $\boldsymbol{s}^*$ is a feasible strategy of the original game. We will try to show that $\boldsymbol{s}^*$ is a strategic Nash equilibrium for the original game. Note also that we can take a subsequence of the $\boldsymbol{s}^*_\epsilon$, say $\{\epsilon_1, \epsilon_2, ...\}$ so that each of the corresponding allocations and prices $\boldsymbol{x}^*_{\epsilon_j}$ and $\boldsymbol{p}^*_{\epsilon_j}$ also converge to a limit point, say $\boldsymbol{x}^*$ and $\boldsymbol{p}^*$, respectively, as they also lie on a compact set. Next we show a lower bound on $\boldsymbol{p}^*_{\epsilon_j}$.

Claim 3. *If at least two players have positive utility for good j, then there is some constant $c > 0$ such that for every epsilon game, the strategic equilibrium price $\boldsymbol{p}^*_\epsilon > c$.*

Proof. We argue by contradiction. Let us choose some ϵ and some good j for which two players have positive utility and such that the equilibrium price is $p^*_{\epsilon j} \leq c$. We will define c later. Since there are at least two users who have positive utility from good j, there is at least one user, say user i, who has $u_{ij} > 0$ but who is allocated at most half of good j (i.e. $x^*_{ij} \leq 1/2$ and could in fact be 0). Consider two cases.

Case 1: $s^*_i \geq \frac{m_i}{2}$.

In this case, by (5.3), we must have $p^*_j \geq y_{ij} = s^*_i u_{ij} \geq \frac{m_i u_{ij}}{2}$. Choosing $c < \frac{m_{\min} u_{\min}}{2}$ gives a contradiction.

Case 2: $s_i^* < \frac{m_i}{2}$.

In this case, recall from (5.29) that:

$$\frac{\partial U_i}{\partial s_i} = \sum_k \frac{\tilde{p}_k}{(\tilde{p}_k + s_i u_{ik}) \prod_j (\tilde{p}_j + s_i u_{ij})^{u_{ij}}} \prod_j u_{ij}^{u_{ij}} \tag{5.30}$$

Since we are assuming $x_{ij}^* < 1/2$, we must have that $\tilde{p}_j^* > y_{ij}^* = s_i^* u_{ij}$. Then, as all of the terms of the above sum are positive, we can simply focus on the j-th term to get the following inequality at the equilibrium point:

$$\frac{\partial U_i}{\partial s_i} > \frac{1}{2(2\tilde{p}_j)^{u_{ij}} \prod_{k \neq j} (\tilde{p}_k + s_i u_{ik})^{u_{ik}}} \prod_k u_{ik}^{u_{ik}} \tag{5.31}$$

Now we let $U = \prod_k u_{ik}^{u_{ik}}$ and notice that each term of the product in the denominator is bounded by the total money between all players (which we will call M). Thus, at equilibrium we have:

$$\frac{\partial U_i}{\partial s_i} > \frac{U}{2(2\tilde{p}_j)^{u_{ij}} M^m} \tag{5.32}$$

Thus, by choosing $c < \frac{1}{2}(\frac{U}{2M^m})^{\frac{1}{u_{\max}}}$, we can ensure that $\frac{\partial U_i}{\partial s_i} > 1$. This contradicts the fact that we are at an internal equilibrium of the strategic game. □

By the above claim it is clear that for each epsilon game the prices for each good must be at least c and, thus, in the limit $\boldsymbol{p}^* > c$. From this we will establish that $\boldsymbol{x}^*$ and $\boldsymbol{p}^*$ are in fact valid prices and allocations for the market equilibrium if the players play strategy $\boldsymbol{s}^*$. First, the demands and prices are feasible as, by convergence, we have that $\sum_i x_{ij}^* = 1$ for all j and $\sum_j x_{ij}^* p_j^* = s_i^*$ for all i. It is also clear from the convergence that the allocation $\boldsymbol{x}^*$ must maximize each player's utility amongst all allocations that they can afford. We need only check that if a player has $\boldsymbol{s}_i^* = 0$ that they are allocated no goods which is the only possible discontinuous condition on the game. This follows from the fact that we have guaranteed that $\boldsymbol{p}^* > c > 0$. Thus, $\boldsymbol{x}^* = \boldsymbol{x}(\boldsymbol{s}^*)$ and $\boldsymbol{p}^* = \boldsymbol{p}(\boldsymbol{s}^*)$.

Since the allocations $\boldsymbol{x}_{\epsilon_j}$ of $\boldsymbol{s}_{\epsilon_j}$ converge to the allocation $\boldsymbol{x}^*$ of $\boldsymbol{s}^*$, it must be that, for every $\delta > 0$, there exists some $J > 0$ such that for all $j > J$:

$$\left|T_i(\boldsymbol{s}^*) - T_i(\boldsymbol{s}_{\epsilon_j})\right| < \delta. \tag{5.33}$$

We are now ready to show that $\boldsymbol{s}^*$ is a Nash equilibrium for the strategic game. Suppose that it is not. Then there must be some player i who has a payoff improving allocation. In fact, suppose that instead of playing $\boldsymbol{s}^*$, player i deviated to a new strategy $\hat{s}_i$ with strictly greater payoff. Define $\hat{\boldsymbol{s}} = (\boldsymbol{s}_1^*, ..., \hat{s}_i, ..., \boldsymbol{s}_n^*)$ and $\hat{\boldsymbol{s}}_\epsilon = (s_{\epsilon_1}, ..., \hat{s}_i, ..., s_{\epsilon_n})$ for sufficiently small ϵ. Again, we partition into two cases.

Case 1: $\hat{s}_i = 0$.

If $\hat{s}_i = 0$ then $s_i^* > 0$. Now consider $s^*_{\epsilon_1 i}, s^*_{\epsilon_2 i}, \ldots$ the set of strategies converging to s_i^*. Since these are at Nash equilibrium, each of these strategies has utility more than $m_i - \epsilon$ (which is the minimum utility obtained if player i only put ϵ in the market in the epsilon game). Thus these must converge to a strategy with utility $\geq m_i$. Thus, defecting with $\hat{s}_i = 0$ which gives utility m_i cannot be a utility increasing move.

Case 2: $\hat{s}_i > 0$.

Suppose $T_i(\hat{\boldsymbol{s}}) - T_i(\boldsymbol{s}^*) = \epsilon' > 0$. Then, for sufficiently small ϵ we must have $T_i(\hat{\boldsymbol{s}}_\epsilon) - T_i(\boldsymbol{s}_\epsilon^*) > 0$ by (5.33). This contradicts the fact that $\boldsymbol{s}_\epsilon^*$ is a Nash equilibrium. Thus $\boldsymbol{s}^*$ must be a Nash equilibrium for the strategic game as required. □

5.5.2 Linear Utility Functions

A Nash equilibrium need not exist in a Fisher Game with linear utilities. We show this using the following simple counterexample. Consider a market with two buyers a and b and two goods 1 and 2. Let each player get utility 1 for each good, except that $u_{a2} = 2$. Let the budgets of each player be $m_a = m_b = 4$. Suppose now that each player chooses a strategy $s_a \leq m_a$ and $s_b \leq m_b$. There are four cases.

Case I: $s_a < s_b$.

The market equilibrium in this case is $p_1 = p_2 = \frac{s_a+s_b}{2}$, a taking only good 2 with total utility $U_a = \frac{4s_a}{(s_a+s_b)} + m_a - s_a$, and b taking the full good 1 and the rest of good 2 with utility $\frac{2s_b}{s_a+s_b} + m_a - s_a$. Now U_a is a concave function in s_a, its derivative is $\frac{4s_b}{(s_a+s_b)^2} - 1$, and the s_a value maximizing it must satisfy $4s_b = (s_a + s_b)^2$, hence this must hold in NE. Similarly, for b, we get $2s_a = (s_a + s_b)^2$ in NE. This gives $s_a = 2s_b$, a contradiction to $s_a < s_b$.

Case II: $s_a = s_b = s$.

Now $s = 0$ cannot be NE, because a buyer putting a tiny amount of money on the market could get the utility 3 or 2, resp. If $s > 0$ then the market equilibrium prices are $p_1 = p_2 = s$, a buying the full unit of 2, b buying the full unit of 1. This cannot be a NE, since if b's utility is $1 + m_b - s_b$ then if he puts in a little less money he will still get the full unit of good 1, giving utility 1 (see next case).

Case III: $s_b < s_a \leq 2s_b$.

At the market equilibrium, a only buys 2 and b only buys 1. Hence $p_1 = s_b, p_2 = s_a$. This clearly cannot be a NE: a's utility is $2 + m_a - s_a$, b's utility $1 + m_b - s_b$, i.e. they get the full utility of the corresponding good for infinitesimal money. In particular, a could decrease s_a.

Case IV: $2s_b < s_a$**.**

At the market equilibrium, $p_1 = \frac{s_a+s_b}{3}$ and $p_2 = \frac{2(s_a+s_b)}{3}$. Buyer a takes the full good 2, b spends all his money on 1. So

$$U_a = \frac{3s_a}{s_a + s_b} - s_a, \quad U_b = \frac{3s_b}{s_a + s_b} - s_b \tag{5.34}$$

Then the same way as in Case I, if $0 < 2s_b < s_a < m_a$, then we must have that if it's a NE then $3s_a = 3s_b = (s_a + s_b)^2$. This again contradicts $2s_b < s_a$.

If $s_b = 0$, then a gets all goods with utility $3 + m_a - s_a$, and could get it for less. If $0 < 2s_b < s_a = m_a = 4$, then again we must have $3s_a = (s_a + s_b)^2$ for b to be optimal, giving $s_b = 2\sqrt{3} - 4 < 0$.

5.6 Social Welfare under Best Response Dynamics

Whilst Nash equilibria need not exist in the Fisher Game with linear utilities, we can still obtain a good welfare guarantee in the dynamic setting. Specifically, in the dynamic setting we assume that in every round (time period), each player simultaneously plays a best response to what they observed in the previous round. Dynamics are a natural way to view how a game is played and a well-studied question is whether or not the game dynamics converge to an equilibrium. Regardless of the answer, it is possible to quantify the average social welfare over time of the dynamic process. This method was introduced by Goemans et al in [78] and we show how it can be applied here to bound the *Dynamic Price of Imperfect Competition* - the worst case ratio of the average welfare of states in the dynamic process to the welfare of the Walrasian equilibrium.

For best responses to be well defined in the dynamic Fisher Game, we need the concept of a minimum monetary allocation s_i. Thus we discretize the game by allowing players to submit strategies which are rational numbers of precision up to Φ. This has the added benefit of making the game finite. In the remainder of this section, we prove the following bound on the Dynamic Price of Imperfect Competition.

Theorem 17. *In the dynamic Fisher Game with linear utilities, the Dynamic Price of Imperfect Competition is lower bounded by* $\Omega(1/\log(\frac{M}{\phi}))$ *where* $M = \max_i m_i$.

To prove Theorem 17, we first notice that if a player puts a certain fraction of his budget onto the market, he is guaranteed at least that fraction of his utility in the Walrasian equilibrium.

Lemma 9. *In strategy profile* $\boldsymbol{s}$ *suppose player* i *has played strategy* $s_i > \frac{m_i}{K}$ *for some* K. *Then* $U_i(\boldsymbol{x}_i(\boldsymbol{s})) \geq \frac{\hat{U}_i}{K}$ *where* $\hat{U}_i$ *is that player's utility in the Walrasian equilibrium.*

Proof. Let β_i and β_i^W be the bang-for-buck of player i at the current strategy and at Walrasian equilibrium, then using Lemma 3 we have $U_i(\boldsymbol{x}_i(\boldsymbol{s})) = s_i\beta_i \geq \frac{m_i}{K}\beta_i \geq \frac{m_i}{K}\beta_i^W = \frac{\hat{U}_i}{K}$. □

Next, we will show that if a player is not receiving much utility in the current strategy state, then in his next move he will either dramatically decrease or dramatically increase his allocation of money to the market.

Lemma 10. *Suppose at time t, the players have chosen strategies $\boldsymbol{s}^t$. If for player i, $T_i(\boldsymbol{s}^t) < \frac{\hat{U}_i}{K}$ then $s_i^{t+1} \geq K s_i^t$.*

Proof. Notice that if for his next move, player i were to put in $s_i^{t+1} = m_i$ then he would get utility at least $\hat{U}_i$ (Lemma 4). Thus his best response must lead him to expect at least this amount. Since increasing s_i from s_i^t will only worsen his bang per buck and reduce the savings, the only way to get at least $\hat{U}_i$ is to put in at least K times what he previously did. □

Lemma 11. *Suppose at time t, the players have chosen strategies $\boldsymbol{s}^t$. If for player i, $T_i(\boldsymbol{s}^t) < \frac{m_i}{K}$, then $s_i^{t+1} \leq \frac{s_i^t}{K}$.*

Proof. Since $T_i(\boldsymbol{s}^t) < \frac{m_i}{K}$, player i's bang-for-buck at $\boldsymbol{s}^t$ is less than $\frac{1}{K}$. Notice that if for his next move, he were to put in $s_i^{t+1} = 0$ then he would get total utility at least m_i. Thus his best response must lead him to expect at least this amount. By Lemma 5, the only way he can expect to increase his bang-for-buck to 1 is by decreasing his allocation of money by a factor of at least $\frac{1}{K}$. □

We observe that it is not possible for the conditions of Lemmas 10 and 11 to be satisfied simultaneously for $K > 1$. We are now ready to prove Theorem 17.

Proof of Theorem 17. Let us fix some constant $K > 1$. We will argue that any player i will receive aggregate utility at least $\frac{\max(\hat{U}_i, m_i)}{K}$ in any sequence of $C \cdot \log(\frac{M}{\phi})$ moves, for some constant C. Note that sum of these agregates is at least $O(\sum_i m_i)$, and therefore the utility of sellers is also taken care of with an additional factor of 2.

Let β_i^W be the bang-for-buck that player i achieves in the Walrasian equilibrium, and let β_i^t be her bang-for-buck in round t. From Lemma 3 we have $\beta_i^t \geq \beta_i^W, \ \forall i, \ \forall t$, and $\hat{U}_i = m_i \beta_i^W$ and $T_i(\boldsymbol{s}^t) = U_i(\boldsymbol{s}^t) + m_i - s_i^t = s_i^t \beta_i^t + m_i - s_i^t$. We will consider 4 cases:

Case I: $1 \leq \beta_i^W \leq K$**.** In this case, player i's bang-for-buck will always be at least 1 in each round. Thus $\forall t, \ T_i(\boldsymbol{s}^t) \geq m_i \Rightarrow T_i(\boldsymbol{s}^t) \geq \frac{\hat{U}_i}{K}$ using $\hat{U}_i = m_i \beta_i^W$.

Case II: $\frac{1}{K} \geq \beta_i^W \geq 1$**.** As $\beta_i^t \geq \beta_i^W, \ \forall i$, we have that she will receive at least $\hat{U}_i$ total payoff which is $\beta_i^W m_i \geq \frac{m_i}{K}$.

Case III: $\beta_i^W > K$**.** Since $\beta_i^t \geq \beta_i^W > K$, we will have that $T_i(\boldsymbol{s}^t) \geq m_i, \ \forall t$. So we need only show that at least once in every $C \cdot \log(\frac{M}{\phi})$ moves, player i receives utility at least $\frac{\hat{U}_i}{K}$. We argue by applying Lemma 10. If player i is not receiving the desired utility, then in the next time period she will increase her allocation by a factor of K. Thus within $O(\log(\frac{M}{\phi}))$ time periods either she receives $\frac{\hat{U}_i}{K}$ payoff or she allocates at least $\frac{m_i}{K}$. In the latter case too she will receive $\frac{\hat{U}_i}{K}$ payoff due to Lemma 9.

Case IV: $\beta_i^W < \frac{1}{K}$**.** Since $\beta_i^W < \frac{1}{K}$, we will have that $T_i(\boldsymbol{s}^t) \geq \hat{U}_i, \ \forall t$. So we need only show that at least once in every $O(\log(\frac{M}{\phi}))$ moves, player i receives utility at least $\frac{m_i}{K}$. In this case, we argue by applying Lemma 11. If player i is not receiving the desired utility, then in the next time period she will decrease her allocation by a factor of $\frac{1}{K}$. Thus, in the next time period she will receive a utility of at least $\frac{m_i}{K}$ which is sufficient. □

5.7 Examples of Fisher Games

In this section, we provide some examples of Fisher Game's to demonstrate some of the range that the Price of Imperfect Competition can achieve.

5.7.1 A Fisher Game with Unbounded PoIC

We begin with a Fisher Game with one good where potential gain in welfare at its only NE is unbounded compared to its WE. Since a CES function on one good is essentially a linear function, we show the result for a Fisher Game under a CES utility function.

Theorem 18. *For any $\Delta > 1$, there exists a Fisher Game under linear utility functions with exactly one NE $\boldsymbol{s}^*$, and $\mathcal{W}(\boldsymbol{s}^*) \geq \Delta \mathcal{W}(\boldsymbol{m})$.*

Proof. Consider the following market with one good a and three buyers $1, 2$ and 3. Buyers 1 has $m_1 = 1$ and $u_{1a} = H$. Buyer 2 is identical: $m_2 = 1$ and $u_{2a} = H$. On the other hand the third buyer has $m_3 = 2L - 2$ and $u_{3a} = 1$. Assuming there is one unit of good j then the market equilibrium is $p_a = 2L$ and $\{x_{1a}, x_{2a}, x_{3a}\} = \{\frac{1}{2L}, \frac{1}{2L}, \frac{2L-2}{2L}\}$. This has a total welfare of

$$\mathcal{W}(\boldsymbol{m}) = \left(\frac{1}{2L} \cdot H + \frac{1}{2L} \cdot H + \frac{2L-2}{2L} \cdot 1\right) + 2L < \frac{H}{L} + 2L + 1 \tag{5.35}$$

There is a Nash equilibrium $\{s_1^*, s_2^*, s_3^*\} = \{1, 1, 0\}$ with $p_j^* = 2$ and $\{x_{1j}^*, x_{2j}^*, x_{3j}^*\} = \{\frac{1}{2}, \frac{1}{2}, 0\}$. For high enough values for H and L, this game has no other equilibrium. The total welfare at this equilibrium is

$$\mathcal{W}(\boldsymbol{s}^*) = \left((\frac{1}{2} \cdot H + 0) + (\frac{1}{2} \cdot H + 0) + (0 \cdot 1 + 2L - 2)\right) + 2 = H + 2L \tag{5.36}$$

Thus, for any $\Delta > 1$, we can choose H high enough relative to L so that the welfare ratio between the Nash equilibrium and the market equilibrium is greater than Δ. $\square$

5.7.2 A Fisher Game with PoIC < 1

In this section we will demonstrate an example of the linear Fisher Game where the PoIC is < 1.

Take a four buyer game with two items. There are three units of good 1 and one unit of good 2 ($e_1 = 3, e_2 = 1$). The buyers have $(m_1, m_2, m_3, m_4) = (1, 1, k+1-\delta, \delta)$ where k is large and $\delta < \frac{6k}{(6k+1)^2}$. The utility coefficients are $(u_{11}, u_{12}) = (3, 0)$, $(u_{21}, u_{22}) = (3, 0)$, $(u_{31}, u_{32}) = (6, 6k)$ and $(u_{41}, u_{42}) = (0, 1)$. Thus buyer 3 is the only buyer who values both goods.

The market equilibrium is $(p_1, p_2) = (1, k)$ with $(x_{11}, x_{12}) = (x_{21}, x_{22}) = (1, 0)$, $(x_{31}, x_{32}) = (1, \frac{k-\delta}{k})$ and $(x_{41}, x_{42}) = (0, \frac{\delta}{k})$. Total welfare at the equilibrium is then

$$\begin{aligned}
&\sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{G}} u_{ij}\cdot x_{ij} + \sum_{i\in\mathcal{B}} m_i \\
&= \left(3\cdot 1 + 3\cdot 1 + (6\cdot 1 + 6k\cdot\frac{k-\delta}{k}) + 1\cdot\frac{\delta}{k}\right) + (1 + 1 + (k+1-\delta) + \delta) \\
&= 7k + 15 + \frac{\delta}{k}\cdot(1-6k) \\
&> 7k + 15 - 6\cdot\delta \qquad (5.37)
\end{aligned}$$

On the other hand, we claim $(m_1^*, m_2^*, m_3^*, m_4^*) = (1, 1, \sqrt{6k\cdot\delta} - \delta, \delta)$ is a Nash equilibrium. This gives the allocation $(x_{11}^*, x_{12}^*) = (x_{21}^*, x_{22}^*) = (\frac{3}{2}, 0)$, $(x_{31}^*, x_{32}^*) =$

$(0, \frac{\sqrt{6k\cdot\delta}-\delta}{\sqrt{6k\cdot\delta}})$ and $(x^*_{41}, x^*_{42}) = (0, \frac{\sqrt{\delta}}{\sqrt{6k\cdot\delta}})$. The welfare of the equilibrium is

$$
\begin{aligned}
&\sum_{i\in\mathcal{B}}\sum_{j\in\mathcal{G}} u_{ij}\cdot x^*_{ij} + \sum_{i\in\mathcal{B}} m_i \\
&= \left(3\cdot\frac{3}{2} + 3\cdot\frac{3}{2} + (6\cdot 0 + 6k\cdot\frac{\sqrt{6k\cdot\delta}-\delta}{\sqrt{6k\cdot\delta}}) + 1\cdot\frac{\delta}{\sqrt{6k\cdot\delta}}\right) \\
&\quad + (1 + 1 + (k+1-\delta) + \delta) \\
&= \left(9 + 6k + \frac{\delta}{\sqrt{6k\cdot\delta}}(1-6k)\right) + (3+k) \\
&= 7k + 12 - \sqrt{6k\cdot\delta} + \sqrt{\frac{\delta}{6k}} \\
&< 7k + 12 \qquad (5.38)
\end{aligned}
$$

As δ is small this is lower welfare than at the Market equilibrium. Now we need to confirm this is a Nash equilibrium. Since player 2 is spending 1 and is only interested in good 1 we must have that $p^*_1 \geq \frac{1}{3}$. Now for buyer 3 to purchase both goods we must have $p^*_2 = k\cdot p^*_1$ and hence $p^*_2 \geq \frac{1}{3}\cdot k$. But only buyers 3 and 4 want good 2 and $m^*_3 + m^*_4 = \sqrt{6k\cdot\delta} < \frac{1}{3}\cdot k$. Thus, for the market to clear, buyer 3 will only purchase good 2.

It follows that we can separate the game in two submarkets. The first has buyers 1 and 2 with good 1, and the second has buyers 3 and 4 with good 2.

Consider the first sub-market. Let's show that buyer 1 is making a best response. She is facing $(m^*_2, m^*_3, m^*_4) = (1, \sqrt{6k\cdot\delta}-\delta, \delta)$ and needs to select m^*_1. When buyer 2 spends $y \leq 1$ dollars, the utility of buyer 1 is $(1-x) + 3\cdot 3\cdot\frac{x}{x+y}$ when she spends $x \leq 1$ dollars. To see this, she wins a $\frac{x}{x+y}$ fraction of the good; there are three units of the good and she gets a utility of 3 per unit.

Taking the derivative we get

$$\begin{aligned} -1 + \frac{9}{x+y} - \frac{9x}{(x+y)^2} &= -1 + \frac{9(x+y) - 9x}{(x+y)^2} \\ &= -1 + \frac{9y}{(x+y)^2} \end{aligned} \quad (5.39)$$

But this is positive because $y = 1$ and $x \leq 1$. Thus buyer will spend as much as possible, that is $x = 1$ is a best response. By symmetry, buyer 2 is also making a best response.

Now consider the second sub-market. When buyer 4 spends y dollars, the utility of buyer 3 is $(k + 1 - \delta - x) + 6k \cdot \frac{x}{x+y}$ when she spends x dollars. To optimise x we equate

$$-1 + \frac{6k}{x+y} - \frac{6kx}{(x+y)^2} = 0 \quad (5.40)$$

$$\therefore \quad 6ky = (x+y)^2 \quad (5.41)$$

$$\therefore \quad \sqrt{6ky} - y = x \quad (5.42)$$

Since buyer 4 is spending δ dollars, it is a best response for buyer 3 to spend $\sqrt{6k \cdot \delta} - \delta$ dollars, as desired.

Now consider buyer 4. When buyer 3 spends x dollars, the utility of buyer 4 is $(1 - y) + 1 \cdot \frac{y}{x+y}$ when she spends $y \leq \delta$ dollars. Taking the derivative we have

$$-1 + \frac{1}{x+y} - \frac{y}{(x+y)^2} = \frac{-(x+y)^2 + x}{(x+y)^2} \quad (5.43)$$

Since buyer 3 is spending $x = \sqrt{6k \cdot \delta} - \delta$ dollars and $y \leq \delta$, the numerator is at least

$$\begin{aligned} x - (x + \delta)^2 &= \sqrt{6k \cdot \delta} - \delta - (\sqrt{6k \cdot \delta})^2 \\ &= \sqrt{6k \cdot \delta} - (6k + 1) \cdot \delta \qquad (5.44) \end{aligned}$$

But this is positive provided

$$6k \cdot \delta > (6k + 1)^2 \cdot \delta^2 \qquad (5.45)$$

$$\therefore \quad \frac{6k}{(6k + 1)^2} > \delta \qquad (5.46)$$

Thus, buyer 4 will spend all his money and we have a Nash equilibrium.

5.8 The Proportional Share Mechanism

In this section we analyze proportional share mechanisms [63] with and without utility for saved money, and compare welfare at corresponding equilibrium. We show that in proportional share mechanisms [63] adding utility for saved money may lead to an unbounded loss in welfare. In other words, the Price of Imperfect Competition may go to zero. This is unlike the Fisher Game, where the Price of Imperfect Competition is bounded below by $\frac{1}{2}$ (Theorem 14). In proportional share mechanisms [63] buyer i allocates in advance a specific amount m_{ij} of money to each good j. The key point here is that when we allow unit utility from each unit of saved money, then prices can rise for some goods.

For example. Take three players and two goods. Let the players have budgets $K, K, 1$, respectively. Let $(u_{11}, u_{12}) = (h^{-1}, 0)$, $(u_{21}, u_{22}) = (h^2, h)$, $(u_{31}, u_{32}) = (0, h^3)$, for some large h.

The optimality conditions at an equilibrium in these games are:

$$u_{ij} \cdot \frac{p_j - m_{ij}}{(p_j^*)^2} = 1 + \epsilon_i^* \qquad \text{if } m_{ij} > 0 \tag{5.47}$$

and

$$u_{ij} \cdot \frac{p_j - m_{ij}}{(p_j^*)^2} \leq 1 + \epsilon_i^* \qquad \text{if } m_{ij} = 0 \tag{5.48}$$

Without having any value for saved money, we have that buyer 1 allocates all her money to good 1 and buyer 3 allocates all his money to good 2. Thus the optimality conditions state if buyer 2 allocates money to both goods then

$$u_{21} \cdot \frac{K}{(K + m_{21})^2} = u_{22} \cdot \frac{1}{(1 + m_{22})^2} \tag{5.49}$$

$$h^2 \cdot \frac{K}{(K + m_{21})^2} = h \cdot \frac{1}{(1 + K - m_{21})^2} \tag{5.50}$$

$$h \cdot (1 + K - m_{21})^2 \cdot K = (K + m_{21})^2 \tag{5.51}$$

But for $h >> K$ this cannot happen and buyer 2 will allocate all her money to good 1. Thus buyer 3 will win all of good 3 fetching social welfare of at least h^3.

On the other hand if each unit of saved money gives unit utility, then buyer 1 will not allocate any money to good 1 unless its price is at most h^{-1}.

Thus player 2 cannot allocate more than h^{-1} to good 1. Thus he allocates at least $K - h^{-1}$ dollars to good 2. Thus the price of good 2 rises! In which case, buyer 3 gets a $\frac{1}{K}$ fraction of good 2. This gives a social welfare of around $\frac{1}{K} \cdot h^3$.

CHAPTER 6
Conclusion

In this thesis, we have made substantial contributions to the application of algorithmic game theory techniques to the economic theory of oligopoly. In particular, we have:

- Formulated the War Chest Minimization Problem for multimarket predatory pricing.
- Demonstrated the NP-hardness of this problem and of multiplicative approximation guarantees.
- Solved this problem in polynomial time in the absence of fixed costs. Also achieved additive approximation guarantees for the Bertrand and Stackelberg models.
- Demonstrated that lookahead search improves the social utility in the Cournot oligopoly context.
- Shown that for generalised second price auctions, lookahead search achieves optimal outcomes under the worst case model and within a constant factor of optimal for the average case model.
- Formulated the Price of Imperfect Competition as a welfare ratio for market games.

- Demonstrated that the Price of Imperfect Competition is at least 2 for the Fisher Game and this can be improved as the level of market competition increases.
- Proved a logarithmic lower bound on the Price of Imperfect Competition in the dynamic Fisher Game.

However, many interesting unanswered questions remain. We outline some of the most promising here.

Additive approximation guarantee for the multimarket Cournot model: While we were able to demonstrate a polynomial time algorithm for finding an additive approximation to the optimal solution of the War Chest Minimization Problem for Stackelberg and Bertrand models, our methodology did not extend directly to the Cournot model. An alternative rounding scheme could be developed to extend our results to this model.

Alternative rules to test for predation: The Areeda-Turner rule is not the only rule used in practice to test for predatory pricing. Adapting our models to other rules could yield a more robust set of results around whether multimarket predation is possible.

What are other contexts for the lookahead model? In this thesis and in an associated paper [116], we explore the impact that using the lookahead search heuristic can have on a large class of games including market games, AdWord auctions, congestion games, network sharing games, and valid and basic utility games. To what other classes of games does the lookahead equilibrium apply? Can one categorize the classes for which this equilibrium is an improvement over the myopic one?

Under what conditions are lookahead equilibria guaranteed to exist? Is there an analogue to Nash's theorem or Rosen's theorem for lookahead equilibria which describes the necessary and sufficient conditions for it to exist?

How are games affected by combining lookahead equilibria with other heuristic search techniques? Our lookahead model explores the simplest of game search heuristics. In practice, this heuristic is often improved through the use of iterative-deepening, branch and bound methods, and imbalanced search trees. Can we model the impact of these advanced algorithmic techniques on game equilibria?

Is the factor 2 bound on the PoIC tight? We proved that for CES and Cobb-Douglas utilities that the Price of Imperfect Competition is at most 2. Is this factor tight or could it be improved by a more subtle analysis of price dynamics?

Bounding the Price of Imperfect Competition for more utility functions: In our analysis, we were able to prove bounds on the Price of Imperfect Competition for Fisher Games where utility functions treated goods like substitutes. The structure of market interactions makes the case of complementary goods (e.g. when the utility functions are Leontieff or CES with $\rho < 0$) more difficult to study. Is a constant Price of Imperfect Competition achievable in this case?

Extending the analysis: The Fisher Game we present is a specific case of the more general Cournot-Walras model introduced by Codognato, Gabszewicz, and Michel. The concept of Price of Imperfect Competition is easy to extend to this model. What bounds are achievable in this context?

REFERENCES

[1] D. Abreu, D. Pearce and E. Stacchetti, "Optimal cartel equilibria with imperfect monitoring", *Journal of Economic Theory*, **39**, pp251-69, 1986.

[2] H. Ackermann, H. Röglin, and B. Vöcking, "On the impact of combinatorial structure on congestion games", *Journal of the ACM*, **55(6)**, 2008.

[3] B. Adsul, C. Babu, J. Garg, R. Mehta, and M. Sohoni, "Nash equilibria in Fisher market", *Proceedings of the 3rd International Symposium on Algorithmic Game Theory (SAGT)*, pp464–475, 2010.

[4] B. Allaz and J.-L. Villa, "Cournot competition, forward markets and efficiency", *Journal of Economic Theory*, **59**, pp1-16, 1993.

[5] P. Areeda and D. Turner, "Predatory pricing and related issues under Section 2 of the Sherman Act", *Harvard Law Review*, **88**, pp997-733, 1975.

[6] K. Arrow, H. Block, and L. Hurwicz, "On the stability of the competitive equilibrium, II", *Econometrica*, **27**, pp82-109, 1959.

[7] K. Arrow and G. Debreu, "Existence of equilibrium for a competitive economy", *Econometrica*, **22**, pp265-290, 1954.

[8] R. Aumann, "Subjectivity and correlation in randomized strategies", *Journal of Mathematical Economics*, **1**, pp67-96, 1974.

[9] B. Awerbuch, Y. Azar and A. Epstein, "The price of routing unsplittable flow", *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, 2005.

[10] B. Awerbuch, Y. Azar, A. Epstein, V. Mirrokni, and A. Skopalik, "Fast convergence to nearly-optimal solutions in potential games", *Proceedings of the 9th ACM Conference on Electronic Commerce (EC)*, pp264–273, 2008.

[11] B. Awerbuch and R. Kleinberg, "Online linear optimization and adaptive routing", *Journal of Computer and System Sciences*, **74(1)**, pp97-114, 2008.

[12] M. Babaioff, B. Lucier, N. Nisan, and R. Paes Leme, "On the efficiency of the Walrasian mechanism", *Proceedings of the 15th ACM conference on Economics and Computation (EC)*, pp783-800, 2014.

[13] M. Balcan, A. Blum, H. Chan, M. Hajiaghayi, "A theory of loss-leaders: making money by pricing below cost", *Proceedings of 3rd International Workshop on Internet and Network Economics (WINE)*, **LNCS 4858**, pp293-299, 2007.

[14] K. Basu and J. Weibull, "Strategy Subsets Closed Under Rational Behaviour", Papers 479, Stockholm - International Economic Studies.

[15] C. Bayer, "The other side of limited liability: predatory behavior and investment timing", *working paper*, 2004.

[16] J. Benoit, "Financially constrained entry into a game with incomplete information", *RAND Journal of Economics*, **15**, pp490-499, 1984.

[17] P. Berenbrink, T. Friedetzky, L. Goldberg, P. Goldberg, Z. Hu, and R. Martin, "Distributed selfish load balancing", *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp354-363, 2006.

[18] J. Bertrand, "Théorie mathématique de la richesse sociale", *Journal des Savants*, pp499-508, 1883.

[19] V. Bilo and M. Flammini, "Extending the notion of rationality of selfish agents: Second Order Nash equilibria", *Theoretical Computer Science*, **412(22)**, pp2296-2311, 2011.

[20] R. Blair and J. Harrison, "Airline price wars: competition or predation", *Antitrust Bulletin*, **44(2)**, pp489-518, 1999.

[21] A. Blum, E. Even-Dar, and K. Ligett, "Routing without regret: on convergence to Nash equilibria of regret-minimizing algorithms in routing games", *Proceedings of the 21st Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pp45-52, 2006.

[22] A. Blum, M. Hajiaghayi, K. Ligett and A. Roth, "Regret minimization and the price of total anarchy", *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pp373-382, 2008.

[23] A. Blum and Y. Mansour, "Learning, regret minimization and correlated equilibria", in *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, E. Tardos, V. Vazirani (eds.), pp79-102, Cambridge University Press, 2007.

[24] J. Bonnisseau and M. Florig, "Existence and optimality of oligopoly equilibria in linear exchange economies", *Economic Theory*, **22(4)**, pp727–741, 2002.

[25] D. Boudreaux, "Predatory pricing in the retail trade: the Wal-Mart case", in *The Economics of the Antitrust Process*, Klewer, pp195-215, 1989.

[26] W. Brainard and H. Scarf, "How to compute equilibrium prices in 1891", *Cowles Foundation Discussion Papers*, No 1272, Cowles Foundation for Research in Economics, Yale University, 2000.

[27] J. Brander, "Export subsidies and international market share rivalry", *Journal of International Economics*, **18**, pp83-100, 1985.

[28] J. Brander and A. Zhange, "Market Conduct in the Airline Industry: An Empirical Investigation", *The RAND Journal of Economics*, **21**, pp567-583, 1990.

[29] S. Branzei, Y. Chen, X. Deng, A. Filos-Ratsikas, S. Frederiksen and J. Zhang, "The Fisher market game: equilibrium and welfare", *Twenty-Eighth Conference on Artificial Intelligence* , 2014.

[30] G. Brown, "Iterative solutions of games by fictitious play", in *Activity Analysis of Production and Allocation*, T. Koopmans (ed.), pp374-376, Wiley, 1951.

[31] T. Bu, X. Deng and Q. Qi, "Forward looking Nash equilibrium for keyword auction", *Information Processing Letters*, **105(2)**, pp41-46, 2008.

[32] M. Burns, "Predatory pricing and the acquisition cost of competitors", *Journal of Political Economy*, **94**, pp266-96, 1986.

[33] M. Busse, "Firm financial condition and airline price wars", *RAND Journal of Economics*, **33(2)**, pp298-318, 2002.

[34] M. Cary, A. Das, B. Edelman, I. Giotis, K. Heimerl, A. Karlin, C. Mathieu and M. Schwarz, "Greedy bidding strategies for keyword auctions", *Proceedings of the ACM International Conference on Electronic Commerce (EC)*, 2007.

[35] E. Chamberlin, *The Theory of Monopolistic Competition*, Harvard University Press, 1932.

[36] X. Chen, D. Dai, Y. Du, and S. Teng, "Settling the complexity of Arrow-Debreu equilibria in markets with additively separable utilities", *50th Annual Symposium on Foundations of Computer Science (FOCS)*, 2009.

[37] N. Chen, X. Deng, H. Zhang, and J. Zhang, "Incentive ratios of Fisher markets", *Proceedings of the 39th International Colloquium on Automata, Languages and Programming (ICALP)*, pp464–475, 2012.

[38] X. Chen and S. Teng, "Spending is not easier than trading: on the computational equivalence of Fisher and Arrow-Debreu equilibria", *Algorithms and Computation*, pp647-656, 2009.

[39] S. Chien and A. Sinclair, "Convergence to approximate Nash equilibria in congestion games", *Games and Economic Behavior*, **71(2)** pp315-327, 2011

[40] G. Christodolou, V. Mirrokni, and A. Sidiropoulos, "Convergence and approximation in potential games", *Proceedings of the 18th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, pp349-360, 2006.

[41] G. Codognato and J. Gabszewicz, "Equilibre de Cournot-Walras dans une économie d'échange", *Revue économique*, **42(6)**, pp1013–1026, 1991.

[42] J. Conlisk, "Why Bounded Rationality?", *Journal of Economic Literature*, **34(2)**, pp669-700, 1996.

[43] A. Cournot, *Recherces sur les Principes Mathématiques de la Théorie des Richesse*, Paris, 1838.

[44] J. Dalton and L. Esposito, "Predatory price cutting and Standard Oil: a re-examination of the trial record", *Research in Law and Economics*, **22**, pp155-205, 2007.

[45] N. Devanur, J. Garg, and L. Végh, "A rational convex program for linear Arrow-Debreu markets", 2013.

[46] N. Devanur, C. Papadimitriou, A. Saberi, and V. Vazirani, "Market equilibrium via a primal-dual algorithm for a convex program", *Journal of the ACM*, **55(5)**, Article 22, 2008.

[47] N. Devanur, C. Papadimitriou, A. Saberi, and V. Vazirani, "Market equilibrium via a primal-dual-type algorithm", *Proceedings of the 43rd Symposium on Foundations of Computer Science (FOCS)*, 2002.

[48] B. Eaves, "A finite algorithm for the linear exchange model", *Journal of Mathematical Economics*, **3(2)**, pp197–203, 1976.

[49] B. Eaves, "Finite solution of pure trade markets with Cobb-Douglas utilities", *Economic Equilibrium: Model Formulation and Solution*, pp226–239, 1985.

[50] B. Edelman, M. Ostrovsky and M. Schwarz, "Internet advertising and the generalised second-price auction: selling billions of dollars worth of keywords", *American Economic Review*, **97(1)**, pp242-259, 2007.

[51] F. Edgeworth, "Me teoria pura del monopolio", *Giornale degli Economisti*, **40**, pp13-31, 1897.

[52] E. Eisenberg and D. Gale, "Consensus of subjective probabilities: The pari-mutuel method", *The Annals of Mathematical Statistics*, **30(1)**, pp165–168, 1959.

[53] E. Eisenberg, "Aggregation of utility functions", *Management Sciences*, **7(4)**, pp337–350, 1961.

[54] G. Ellison, "Learning, Local Interaction, and Coordination", *Econometrica*, **61**, pp1047-1071, 1993.

[55] G. Ellison, "Theories of cartel stability and the joint executive committee", *RAND Journal of Economics*, **25**, pp37-57, 1994.

[56] K. Elzinga and D. Mills, "Testing for predation: Is recoupment feasible?", *Antitrust Bulletin*, **34**, pp869-893, 1989.

[57] K. Elzinga and D. Mills, "Price wars triggered by entry", *International Journal of Industrial Organization*, **17**, pp179-198, 1999.

[58] R. Ericson and A. Pakes, "Markov-perfect industry dynamics: A framework for empirical work", *The Review of Economic Studies*, **62(1)**, pp53-82, 1995.

[59] E. Even-Dar, Y. Mansour, and U. Nadav, "On the convergence of regret minimization dynamics in concave games", *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, 2009.

[60] A. Fabrikant and C. Papadimitriou, "The complexity of game dynamics: BGP oscillations, sink equlibria, and beyond", *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp844-853, 2008.

[61] A. Fabrikant, C. Papadimitriou, and K. Talwar, "The complexity of pure Nash equilibria", *Proceedings of the 36th Annual ACM Symposium on Theory of Computing (STOC)*, pp604-612, 2004.

[62] A. Fanelli, M. Flammini, and L. Moscardelli, "The speed of convergence in congestion games under best-response dynamics", *Proc. of the 35th International Colloquium on Automata, Languages and Programming (ICALP)*, pp796-807, 2008.

[63] M. Feldman, K. Lei, and L. Zhang, "The proportional-share allocation market for computational resources", *Proceedings of the 6th ACM Conference on Electronic Commerce (EC)*, 2005.

[64] C. Fershtman and A. Pakes, "A dynamic oligopoly with collusion and price wars", *RAND Journal of Economics*, **31(2)**, pp207-236, 2000.

[65] F. Fisher, "The IBM and Microsoft cases: what's the difference?", *American Economic Review*, **90(2)**, pp180-183, 2000.

[66] L. Fortnow and R. Santhanam, "Bounding rationality by discounting time", *Proceedings of The First Symposium on Innovations in Computer Science (ICS)*, 2010.

[67] D. Foster and R. Vohra, "Calibrated learning and correlated equilibrium", *Games and Economic Behavior*, **21**, pp40-55, 1997.

[68] M. Friedman, "The methodology of positive economics", in *Essays in Positive Economics*, M. Friedman, University of Chicago Press, pp3-43, 1953.

[69] D. Fudenberg and J. Tirole, "A 'signal-jamming' theory of predation", *RAND Journal of Economics*, **17**, pp366-376, 1986.

[70] J. Gabszewicz and P. Michel, "Oligopoly equilibria in exchange economies", *Trade, technology and economics. Essays in honour of Richard G. Lipsey*, pp217–240, 1997.

[71] D. Gale, "The linear exchange model", *Journal of Mathematical Economics*, **3(2)**, pp205-209, 1976.

[72] R. Garg and S. Kapoor, "Auction algorithms for market equilibrium", *Proceedings of the 36th Symposium on the Theory of Computing (STOC)*, 2004.

[73] D. Genesove and W. Mullin, "Predation and its rate of return: the sugar industry 1887-1914", *RAND Journal of Economics*, **37(1)**, pp47-69, 2006.

[74] G.V. Gens and E.V. Levner, "Computational complexity of approximation algorithms for combinatorial problems", in *Mathematical Foundations of Computer Science* (*Lecture Notes in Computer Science*, **74**), Springer, Berlin, pp292-300, 1979.

[75] R. Gibbons, *A Primer in Game Theory*, Harvester Wheatsheaf, 1992.

[76] G. Gigerenzer and R. Selten (eds), *Bounded Rationality: the Adaptive Toolbox*, MIT Press, 2001.

[77] M. Goemans, L. Li, V. Mirrokni, and M. Thottan, "Market sharing games applied to content distribution in ad-hoc networks", *Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing (MobiHoc)*, 2004.

[78] M. Goemans, V. Mirrokni, and A. Vetta, "Sink equilibria and convergence", *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*, 2005.

[79] E. Green and R. Porter, "Noncooperative collusion under imperfect price competition", *Econometrica*, **52**, pp87-100, 1984.

[80] A. de Groot, *Thought and Choice in Chess*, 2nd Edition, Mouton, 1978. [Original Version: *Het denken van den Schaker, een experimenteel-psychologische studie*, Ph.D. thesis, University of Amsterdam, 1946.]

[81] M. Güntzer and D. Jungnickel, "Approximate minimization algorithms for the 0/1 Knapsack and Subset-Sum Problem", *Operations Research Letters*, **26(2)**, pp55-66, 2000.

[82] S. Hart and A. Mas-Colell, "A simple adaptive procedure leading to correlated equilibrium", *Econometrica*, **68(5)**, pp1127-1150, 2000.

[83] S. Hart "Adaptive Heuristics", *Econometrica*, **73(5)**, pp1401-1430, 2005.

[84] H. Hotelling, "Stability in Competition", *Economic Journal*, **41**, pp41-57, 1929.

[85] L. Hurwicz, "On informationally decentralized systems", In *Decision and Organization: A volume in Honor of Jacob Marschak*, Volume 12 of Studies in Mathematical and Managerial Economics, pp297–336, 1972.

[86] R. Isaacs, *Differential game: a mathematical theory with applications to warfare and pursuit, control and optimization*, John Wiley and Sons, 1965.

[87] K. Jain, "A polynomial time algorithm for computing an Arrow-Debreu equilibrium for linear utilities", *45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pp286–294, 2004.

[88] K. Jain, "A polynomial time algorithm for computing an Arrow-Debreu equilibrium for linear utilities", *SIAM Journal on Computing*, **37(1)**, pp291–300, 2007.

[89] K. Jain and V. Vazirani, "Eisenberg–Gale markets: algorithms and game-theoretic properties", *Games and Economic Behavior*, **70(1)**, pp84–106, 2010.

[90] P. Jehiel, "Limited horizon forecast in repeated alternate games", *Journal of Economic Theory*, **67**, pp497-519, 1995.

[91] R. Johari and J. Tsitsiklis, "Efficiency loss in network resource allocation game", *Mathematics of Operations Research*, **57(4)**, pp823-839, 2004.

[92] D. Kahneman, "Maps of bounded rationality: psychology for behavioral economics", *The American Economic Review*, **93(5)**, pp1449-1475, 2003.

[93] D. Kahneman, P. Slovic and A. Tversky (eds), *Judgement under Uncertainty: Heuristics and Biases*, pp201-208, Cambridge University Press, 1982.

[94] D. Kahneman and A. Tversky, "The simulation heursitic", in *Judgement under Uncertainty: Heuristics and Biases*, D. Kahneman, P. Slovic and A. Tversky (eds), pp201-208, Cambridge University Press, 1982.

[95] M. Kandori, G. Mailath and R. Rob, "Learning, mutuation, and long-run equilibria in games", *Econometrica*, **61**, pp29-56, 1993.

[96] F. Kelly and V. Vazirani, "Rate control as a market equilibrium", Unpublished manuscript, 2002.

[97] J. de Kleer and O. Raiman, "How to diagnose with very little information", *Fourth International Workshop on Principles of Diagnosis*, pp160-165, 1993.

[98] J. de Kleer, O. Raiman and M. Shirley, "One step lookahead is pretty good", in *Readings in Model-Based Diagnosis*, W. Hamscher, L. Console, and J. de Kleer (eds), pp138-142, Morgan Kaufmann, 1992.

[99] P. Klemperer, "Price wars caused by shifting costs", *Review of Economic Studies*, **56**, pp405-420, 1989.

[100] G. Klein, "Developing expertise in decision making", *Thinking and Reasoning*, **3(4)**, pp337-352, 1997.

[101] G. Klein, *Sources of Power: How People make Decisions*, MIT Press, 1998.

[102] B. Klein, "The Microsoft case: what can a dominant firm do to defend its market position", *Journal of Economic Perspectives*, **15(2)**, pp45-62, 2001.

[103] B. Klein, "Did Microsoft engage in anticompetitive exclusionary behavior" *Antitrust Bulletin*, **46(1)**, pp71-113, 2001.

[104] R. Kleinberg, G. Piliouras and E. Tardos, "Multiplicative updates outperform generic no-regret learning in congestion games", *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, 2009.

[105] E. Kohlberg and J. Mertens, "On the strategic stability of equilibria", *Econometrica*, **54(5)**, pp1003-1037, 1986.

[106] E. Koutsoupias and C. Papadimitriou, "Worst-case equilibria", *STACS*, 1999.

[107] D. Kreps and J. Scheinkman, "Quantity precommitment and Bertrand competition yield Cournot outcomes", *The Bell Journal of Economics*, **14**, pp326-337, 1983.

[108] R. Paes Leme and E. Tardos, "Pure and Bayes-Nash price of anarchy for generalized second price auction", *Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, 2010.

[109] A. Lerner and H. Singer, "Some notes on duopoly and spatial competition", *Journal of Political Economy*, pp145-186, 1937.

[110] B. Lipman, "Information processing and bounded rationality: a survey", *The Canadian Journal of Economics*, **28(1)**, pp42-67, 1995.

[111] E. Markakis and O. Telelis, "Discrete strategies in keyword auctions and their inefficiency for locally aware bidders", *Proceedings of the 6th Workshop on Internet and Network Economics (WINE)*, 2010.

[112] J. McGee, "Predatory price cutting: the Standard Oil (NJ) case", *Journal of Law and Economics*, **1**, pp137-169, October, 1958.

[113] R. Mehta, N. Thain, László Végh and A. Vetta, "To save or not to save: the Fisher game", *Proceedings of 11th International Workshop on Internet and Network Economics (WINE)*, pp294-307, 2014.

[114] P. Milgrom and J. Roberts, "Predation, reputation, and entry deterrence", *Journal of Economic Theory*, **27**, pp280-312, 1982.

[115] V. Mirrokni and A. Skopalik, "On the complexity of Nash dynamics and sink equilibria", *Proceedings of the ACM International Conference on Electronic Commerce (EC)*, 2009.

[116] V. Mirrokni, N. Thain and A. Vetta, "A theoretical examination of practical game playing: lookahead search", *Proceedings of the 5th international conference on Algorithmic Game Theory (SAGT)*, pp251-262, 2012.

[117] V. Mirrokni and A. Vetta, "Convergence issues in competitive games", *Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pp183-194, 2004.

[118] A. Montanari and A. Saberi, *Convergence to equilibrium in local interaction games and ising models*, Technical Report, arXiv:0812.0198, CoRR, 2008.

[119] R. Myerson, "Nash Equilibrium and the History of Economic Theory", *Journal of Economic Literature*, **37(3)**, pp1067-1082, 1999.

[120] R. Nagel, "Unraveling in guessing games: an experimental study", *The American Economic Review*, **85(5)**, pp1313-1326, 1995.

[121] D. Nau, "Pathology on game trees: a summary of results", *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, pp102-104, 1980.

[122] D. Nau, "An investigation of the causes of pathology in games", *Artificial Intelligence*, **19**, 257-278, 1982.

[123] D. Nau, "Decision quality as a function of search depth on game trees", *Journal of the ACM*, **30(4)**, pp687-708, 1983.

[124] D. Nau, "Pathology on game trees revisited, and an alternative to minimaxing", *Artificial Intelligence*, **21**, pp221-244, 1983.

[125] A. Newell and H. Simon, *Human Problem Solving*, Prentice-Hall, 1972.

[126] N. Nisan, J. Bayer, D. Chandra, et al., "Google's auction for TV ads", *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP)*, pp309-327, 2009.

[127] P. Oberender and T. Rudolf, "Heinrich Von Stackelberg (1905-1946)", In *Pioneers of Industrial Organization*, pp50-51, 2007.

[128] J. Ordover and G. Saloner, "Predation, monopolization, and anti-trust", in R. Schmalensee and R. Willig (eds.), *The Handbook of Industrial Organization*, North-Holland, pp537-596, 1989.

[129] C. Papadimitriou, "On the complexity of the parity argument and other inefficient proofs of existence", *Journal of Computer and System Sciences*, **48**, pp498-532, 1994.

[130] C. Papadimitriou, "Algorithms, games, and the internet", *Proceedings of the 33rd annual ACM Symposium on Theory of Computing (STOC)*, pp749-753, 2001.

[131] C. Papadimitriou, "The Complexity of Finding Nash Equilibria", in *Algorithmic Game Theory*, pp29-50, 2007.

[132] J. Pearl, *Heuristics: Intelligent Search Strategies for Computer Problem Solving*, Addison-Wesley, 1984.

[133] J. Podolny and F. Scott Morton, "Social status, entry and predation: the case of British shipping cartels 1879-1929", *Journal of Industrial Economics*, **47(1)**, pp41-67, 1999.

[134] R. Porter, "On the incidence and duration of price wars", *The Journal of Industrial Economics*, **33(4)**, pp415-426, 1985.

[135] T. Puu and I. Sushko, *Oligopoly dynamics: Models and tools*, Springer, 2002.

[136] A. Rao, M. Bergen and S. Davis, "How to fight a price war", *Harvard Business Review*, pp107-120, March/April, 2000.

[137] J. B. Rosen, "Existence and Uniqueness of Equilibrium Points for Concave N-person Games", *Econometrica*, **33(3)**, pp520-534, 1965.

[138] J. Rotemberg and G. Soloner, "A supergame-theoretic model of price wars during booms", *American Economic Review*, **76(3)**, pp390-407, 1986.

[139] A. Rubenstein, *Modeling Bounded Rationality*, MIT Press, 1998.

[140] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 2nd Edition, Prentice-Hall, 2002.

[141] L. Savage, *The Foundation of Statistics*, Wiley, 1954.

[142] T. Sargent, *Bounded Rationality in Macroeconomics*, Clarendon Press, 1993.

[143] H. Scarf, "Some Examples of Global Instability of the Competitive Equilibrium", *Cowles Foundation Discussion Papers*, **79**, 1959.

[144] H. Scarf, "The Computation of Equilibrium Prices", in *Applied General Equilibrium Analysis*, Cambridge Books, 1984.

[145] F. Scherer, *Industrial Market Structure and Economic Performance*, Houghton Miffin, 1980.

[146] F. Scott Morton, "Entry and predation: British shipping cartels 1879-1929", *Journal of Economics and Management Strategy*, **6**, pp679-724, 1997.

[147] E. Sefer and U. Kuter and D. Nau, "Real-time A* Search with Depth-k Lookahead", *Proceedings of the International Symposium on Combinatorial Search*, 2009.

[148] R. Selten, "The chain store paradox", *Theory and Decision*, **9**, pp127-159, 1978.

[149] R. Selten, "What is bounded rationality?", in *Bounded Rationality: the Adaptive Toolbox*, G. Gigerenzer and R. Selten (eds), MIT Press, pp13-36, 2001.

[150] R. Selten, "Boundedly Rational Qualitative Reasoning on Comparative Statics", in *Advances in Understanding Strategic Behavior: Game Theory, Experiments and Bounded Rationality*, Steffen Huck (ed.), Palgrave Macmillan, pp1-8, 2004.

[151] A. Sen, "Rational fools: a critique of the behavioral foundations of economic theory", *Philosophy and Public Affairs*, **6(4)**, pp317-344, 1977.

[152] A. Sen, "Internal consistency of choice", *Econometrica*, **61(3)**, pp495-521, 1993.

[153] C. Shannon, "Programming a computer for playing chess", *Philosophical Magazine*, Series 7, **41(314)**, pp256-275, 1950.

[154] L. Shapley and M. Shubik, "Trade using one commodity as a means of payment", *The Journal of Political Economy*, pp937–968, 1977.

[155] V. Shmyrev, "An algorithm for finding equilibrium in the linear exchange model with fixed budgets", *SIAM Journal of Applied and Industrial Mathematics*, **3(4)**, pp505–518, 2009.

[156] H. Simon, "A behavioral model of rational choice", *Psychological Review*, **63**, pp129-138, 1955.

[157] H. Simon, "Rational choice and the structure of the environment", *Psychological Review*, **63**, pp129-138, 1956.

[158] H. Simon, *The Sciences of the Artificial*, 3rd edition, MIT Press, 1996.

[159] A. Skopalik and B. Vöcking, "Inapproximability of pure Nash equilibria", *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC)*, pp355-364, 2008.

[160] M. Slade, "Strategic pricing models and interpretation of price war data", *European Economic Review*, **34**, pp524-537, 1990.

[161] R. Solov, "A contribution to the theory of economic growth", *Quarterly Journal of Economics*, **70**, pp65–94, 1956.

[162] H. von Stackelberg, *Marktform und Gleichgewicht*, Julius Springer, Vienna, 1934.

[163] D. Stahl and P. Wilson, "Experimental evidence on players' models of other players", *Journal of Economic Behavior and Organization*, **25(3)**, pp309-327, 1994.

[164] D. Stahl and P. Wilson, "On players' models of other players: theory and experimental evidence", *Games and Economic Behavior*, **10(1)**, pp218-254, 1995.

[165] G. Stigler, "A theory of oligopoly", *Journal of Political Economy*, **72(1)**, pp44-61, 1964.

[166] J. Swann, "Stackelberg, multistage games", *6230 Spring 2007: Economic Decision Analysis II*, website: http://www2.isye.gatech.edu/ jswann/teaching/6230S08.htm#Lectures.

[167] N. Thain and A. Vetta, "Computational aspects of multimarket price wars", *Proceedings of 5th International Workshop on Internet and Network Economics (WINE)*, pp304-315, 2009.

[168] J. Tirole, *The Theory of Industrial Organization*, MIT Press, 1988.

[169] L. Telser, "Cutthroat competition and the long purse", *Journal of Law and Economics*, **9**, pp259-277, 1966.

[170] A. Tversky and D. Kahneman, "Judgement under uncertainty: heuristics and biases", *Science*, **185(4157)**, pp1124-1131, 1974.

[171] H. Varian, "Position auctions", *International Journal of Industrial Organization*, **25**, pp1163-1178, 2007.

[172] V. Vazirani, "Combinatorial algorithms for market equilibria", in *Algorithmic Game Theory*, pp103-133, 2007.

[173] V. Vazirani, M. Yannakakis, "Market equilibrium under separable, piecewise-linear, concave utilities", *Innovations in Computer Science*, pp156-165, 2010

[174] A. Vetta, "Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions", *Proceedings of 43rd Symposium on Foundations of Computer Science (FOCS)*, pp416-425, 2002.

[175] L. Walras, "Principe d'une théorie mathématique de l'échange", 1874.

[176] J. Wood, *Joseph A. Schumpeter: Critical Assessments*, Psychology Press, 1991.

[177] G. Weintraub and C.L. Benkard and B.V. Roy, "Markov perfect industry dynamics with many firms", *Econometrica*, **76(6)**, pp1375-1411, 2008.

[178] B. Willems, "Modeling cournot competition in an electricity market with transmission constraints", *The Energy Journal*, **23**, pp95-125, 2002.

[179] Y. Ye, "Computing the Arrow-Debreu competitive market equilibrium and its extensions", *Proceedings of the 1st International Conference on Algorithmic Applications in Management*, pp3-5,2005.

[180] H. Young, "The evolution of conventions", *Econometrica*, **61**, pp57-84, 1993.

[181] H. Young, *Strategic Learning and its Limits*, Oxford University Press, 2004.