ABSTRACT

CAI, GANGSHU. Flexible Decision-Making in Sequential Auctions. (Under the direction of Assistant Professor Peter R. Wurman).

Because sequential auctions have permeated society more than ever, it is desirable for participants to have the optimal strategies beforehand. However, finding closed-form solutions to various sequential auction games is challenging. Current literature provides some answers for specific cases but not for general cases. A decision support system that can automate optimal bids for players in different sequential auction games will be useful in solving these complex economic problems, which requires not only economic but also computational efficiency.

This thesis contributes in several directions. First, this dissertation derives results related to the multiplicity of equilibria in first-price, sealed-bid (FPSB) auctions, and sequential FPSB auctions, with discrete bids under complete information. It also provides theoretical results for FPSB auctions with discrete bids under incomplete information. These results are applicable to both two-person and multi-person cases.

Second, this thesis develops a technique to compute strategies in sequential auctions. It applies Monte Carlo simulation to approximate perfect Bayesian equilibrium for sequential auctions with discrete bids and incomplete information. It also utilizes the leveraged substructure of the game tree which can dramatically reduce the memory and computation time required to solve the game. This approach is applicable to sequences of a wide variety of auctions.

Finally, this thesis analyzes the impact of information in sequential auctions with continuous bids and incomplete information when bids are revealed. It provides theoretical results especially the non-existence of pure-strategy symmetric equilibrium in both the symmetric sequential FPSB and the symmetric sequential Vickrey auctions.

Flexible Decision-Making in Sequential Auctions

by

Gangshu Cai

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial satisfaction of the
requirements for the Degree of
Doctor of Philosophy

Operations Research and Computer Science

Raleigh

2005

Approved By:

Dr. Jon Doyle	Dr. Salah E. Elmaghraby
Dr. Peter R. Wurman Chair of Advisory Committee	Dr. Xiuli Chao

To my wife Yu Cheng

&

To my parents Genchi Cai and Shuiying Xu

Biography

Cai, Gangshu was born in Zhangzhou, China. He received his B.S. in Physics from the Department of Physics in 1996, and then his M.S. in Economics and Applied Statistics from the Guanghua School of Management in 1999, both from Peking University. He studied Marketing in the Business School at Chinese University of Hong Kong for one year before joining the graduate program in Operations Research and Computer Science at North Carolina State University in 2000. He also worked as an intern at IBM Thomas J. Watson Research Center in New York. He is a member of AAAI, ACM, and INFORMS.

Acknowledgements

I would like to express my deepest gratitude and appreciation to my advisor Dr. Peter R. Wurman. This dissertation would not have come into being without his guidance, inspiration, and support. His feedbacks, including those numerous handwritten comments among lines on my manuscripts, have helped shape my research into the right direction. His unmatched technical skills, as well as his peerless kindness, have made him a great academic role model.

The committee plays a significant role in my dissertation. I am very grateful to Dr. Xiuli Chao for his encouragement, instructions, and generosity in our every discussion. He is my de facto co-advisor. His impact on my dissertation is tremendous and will continue to be in the future. I appreciate the supportiveness and kindness of Dr. Salah E. Elmaghraby. I feel greatly privileged to have learned from such renowned, knowledgeable, and admirable professor. I thank Dr. Jon Doyle for his very valuable comments and recommendations for my dissertation. I also thank Dr. Yahya Fathi for his generous participation on the final defense and review of the dissertation.

Five years might be a short period; however, it has become an unexpected lengthy period for both my wife and me. My wife, Yu, started her graduate study at University of Virginia, one year before I came to North Carolina in 2000. After a year of four-hour driving every other week from Raleigh to Charlottesville, Yu graduated in 2001 and started to work in Chicago. She was relocated to Phoenix, Arizona, in 2004. While I am wrapping up the final stage of this dissertation, we have to commute from coast to coast to see each other. I cannot count exactly how many times she was so helpless on the other end of line. I feel short of words to describe her endurance and braveness to combat with the pain of loneliness at those difficult moments in these years. At the same time, I am overwhelmed by her endless love and support. I owe the existence of this dissertation to her, and hopefully, I can repay her at least partially in the rest of my life.

My parents have shown their mostly selfless love, care, support, and confidence in me. I still remember those days when my father rode a bicycle for miles to see me at school. I also clearly remember those most merciful moments when my mother prays for me. They are always the utmost source when I seek for impetus and strength.

I am indebted to Jih-Shyr Yih and Trieu C. Chieu for their generosity to offer me an internship in the e-Commerce Architecture Department at IBM Watson Research Center. This opportunity enabled me to learn a lot from IBM researchers and to explore into a new research domain in e-business.

My appreciation will be extended to my friends at North Carolina State University. Hao

Cheng, Yujun Wu, Hao Zhang, and Jie Zhong provided very valuable suggestions and help for my dissertation. I am indebted to members of Intelligent Commerce Research Group, especially Tiejun li, Ashish Sureka and Weili Zhu for their help and friendship. I would like to thank other friends who have brought me many cherishable memories in Raleigh. I thank these friends, and the many others, for the pleasant times I have spent with them, and I wish them all luck in their future endeavors.

Gangshu

February 7, 2005

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Auctions have permeated into our society more than ever. More and more companies utilize auctions as an important channel in marketing their products. Millions of people purchase commodities from Internet auction sites, such as eBay, Priceline.com, Yahoo Auctions, and Amazon Auctions. As a glimpse of the size of these auction markets, in 2003 the gross revenue of eBay reached 15 billion [34].

What is an auction? An auction is a market institution in which prices and resource allocation are determined by an explicit set of rules on the basis of bids from the market participants [62].

An auction is a dynamic pricing tool that allows sellers and buyers to reach an agreement on prices and allocations. Either sellers or buyers or third parties can initiate an auction. Sellers might want to utilize auctions to sell items at a higher price to some more affordable customers, and hence to increase the overall revenue; while buyers might enjoy a more flexible market when participating in auctions and avoid overpaying. Usually in auctions, the market is efficient when the buyers with the highest valuation win the items.

The research on auctions has burgeoned in the past decades. Many auction mechanisms, especially the standard auctions like the English, the Dutch, the first-price sealed-bid, and the Vickrey auctions [41], have been discussed in great detail in literature (see, for example [41, 51, 62, 72, 73]). The auction mechanisms discussed in this thesis include:

- First-Price Sealed Bid (FPSB) Auction: In the first-price sealed-bid auction, each bidder submits a single bid independently, without observing others' bids, and the winner with the highest bid pays the price of the highest bid.

- Second-Price Sealed Bid (SPSB or *Vickrey*) Auction: In the Vickrey auction, each bidder submits a single bid independently, without observing others' bids, and the winner with the highest bid pays the price of the second highest bid.

- English Auction: In the English auction, the price is successively increased until only one bidder remains, and the winner pays the final price.

- Dutch Auction: In the Dutch auction, the auctioneer starts at a high price, and then lowers the price continuously. The first bidder who calls out wins the object and pays at the current price.

- *Mth-Price* Auction: There are M objects for sale in an *Mth-price* auction. In *Mth-price* auctions, winners pay at the price of the lowest winning bid. The Mth-price is a little bit different from the *uniform* auction [106], in which winners pay the highest rejected bid or $(M+1)th$ price, like in the Vickery auction.

- *Pay-Your-Bid* Auction: A *pay-your-bid* auction is a multi-object auction, in which winners pay the prices they bid. *Pay-your-bid* auctions are also classified as *discriminatory* auctions in some literature [106], because winners pay different prices for identical items.

Wurman, et al. present an auction parametrization which is useful for designing auction mechanisms [112]. In their work, the parameterization of the auction design space is broad enough to encompass most of the classic auctions and many others [109]. There are three axes are introduced: bidding rules, clearing policy, and information revelation policy. The elements in each axis are listed as following.[1]

1. Bidding Rules

   - Restrictions on sellers; buyers; objects; number of auctions; expressiveness; bid refinements; schedule; activity.

2. Clearing Policy

   - Clear timing; closing conditions; matching function; tie breaking; auctioneer fees.

[1] Interested readers please refer to Wurman et al.[112] for more details.

3. Information Revelation Policy

   - Price quotes; quote timing; order book; transaction history.

A sequential auctions is a market scenario that consists of a sequence of individual auctions. In both local auction houses and on-line auction sites it is quite common to see identical or nearly-identical items sold in a sequence. Examples include auctions for electronic devices, art, wine, fish, flowers, mineral rights, satellite broadcast licenses, government debts, and many others [27]. Among those reported in the academic literature are the sequential sale of 120 identical cases of wine in 1990 at Christie's of Chicago [63] and the sale of pelts on the Seattle Fur Exchange [52]. eBay, the world's largest electronic auction, can be viewed as an unending series of auctions for hundreds of thousands of nearly identical items.

The vast number of trading opportunities and the increasingly fluid markets bolsters the need for automated trading support in the form of *trading agents*—software programs that participate in electronic markets on behalf of a user. Simple bidding tools, like eSnipe[2] and AuctionBlitz[3] enable bidders to automate submission of last-second bids on eBay. However, these tools lack the sophistication that bidders require when faced with a plethora of sequential auctions possibly hosted at multiple auction sites.

The literature on sequential auctions dates back to Vickrey [102], in which he obtains an equilibrium solution for a sequence of first-price auctions with bidders whose single-unit-demand valuations are drawn from a uniform distribution. Since Vickrey's original work, a great deal of research has been directed towards understanding sequential auctions. Milgrom and Weber [74] discuss the equilibrium solutions and price trends under more general assumptions. The following year, Weber published a sequential auction model [106] that served as a foundation for many of the papers that followed.

The rest of the literature on sequential auctions identifies a wide variety of research areas. Bernhardt and Scoones [5] find that a more dispersed valuation distribution on one item may yield more revenue for the seller. Gale and Stegeman [27] model two completely informed and asymmetric buyers bidding for N identical objects from N sellers sequentially under complete information by assuming that the value of one object depends on the number of objects obtained. Branco [8] models a two-unit sequential English auction when some bidders have superadditive (complementary) values for the objects. Sørensen [97] finds that, in theory, objects are allocated as a bundle

[2] `http://www.esnipe.com`
[3] `http://www.auctionblitz.com`

more often than as independently in sequential auctions for complements.

A sequence of prices is a *martingale* if prices drift neither up nor down over time [106]. Milgrom and Weber [74, 106] predict a martingale among the price trend in symmetric equilibrium of single-unit demand sequential auctions. However, experiments often show the price declines, which is called the *declining price anomaly* or *afternoon effect* [63, 85]. Beggs and Graddy [4] report some empirical afternoon effect results from art auctions. Some researches find that the anomaly is explained by varying the assumptions. For example, Engelbrecht-Wiggans finds that prices will on average have a downwards trend in a sequence of auctions for a large enough number of stochastically equivalent objects with bounded values [17]. McAfee and Vincent explain that sequential auctions with risk averse bidders will have a decreasing pattern of prices [63]. Gale and Stegeman [27] claim prices decline weakly along any equilibrium path in a multi-unit demand model with two asymmetric buyers. Katzman [39] concludes that the price trend may decrease in expectation in a game of two second price auctions with multiunit demand, symmetric, incomplete information, when there is a high degree of *ex ante* asymmetry of bidder beliefs.

Pitchik and Schotter [85] present some laboratory results from an experiment with budget-constrained, perfectly informed bidders. They conclude that bidders attempt to exploit the constraints of others, and in doing so, bidders might bid up the prices in early stages. As a result, the opponents might deplete their budgets and the later auctions might become less competitive [41, 85].

It is commonly believed that bidders' behaviors will change when they are forced to pay an entry fee or when there is a reserve price. von der Fehr [103] shows that prices will typically decline for later units in a model with participation constraints, e.g., entry fee. McAfee and Vincent [64] prove revenue equivalence between repeated first price and second price sequential auctions with reserve price.

In symmetric, single-unit demand, risk-neutral settings, the revelation of winners and the winning bids in the previous auctions has no effect on the forthcoming auction [106]. Jeitschoko [36] points out that it might be due to the continuous properties of valuation distribution. He also explicitly models an auction where each bidder has only two types, either high valuation or low valuation. In this model, the winner's price information revealed in the first auction has significant influence on the equilibrium bids for both bidders in the second auction. The seminal work by Ortega-Reichert [86] shows a learning process from the signals revealed in the first stage in a two stage model. Hausch [32] generalizes Ortega-Reichert's model and provides necessary conditions for symmetric equilibrium in sequential second-price and first-price auctions. Engelbrecht-Wiggans

[20] finds that in two-bidder, multi-unit demand, sequential auctions, an uninformed bidder may have strictly more expected profit than an informed bidder.

Elmaghraby [15] shows that the order in which heterogeneous items are auctioned will influence the outcome. Gale and Hausch [25] show that giving the buyer the right-to-choose her preferred item from the remaining items induces declining prices. Jeitschko [37] models $n \geq 3$ single-unit demand bidders in a sequential auction with a stochastic number of identical objects. Gale, Hausch and Stegeman [26] model two identical suppliers in sequential second price auctions with subcontracting. Krishna [50] shows that deterring entry at one stage affects the cost of doing so in later stages in a monopolist model.

Recently, the design of more sophisticated trading agents has attracted the attention of researchers in artificial intelligence and other related fields [9, 30, 90, 100, 108]. In most of these studies, the agents are designed for a particular marketplace and lack flexibility to adapt to other market configurations.

The vast majority of auction research models them as games. An equilibrium strategy is a stable solution in which no player wants to unilaterally deviate from the strategy profile. Thus, finding optimal strategies in auctions is naturally transformed to finding the equilibrium strategies in the auction games.

The literature on sequential auctions provides answers for specific cases but not for general cases. When the strategy space is discrete and finite, however, the cost of doing the computation increases exponentially in terms of memory and computation time. When the strategy space is continuous or infinite, to date, there are few explicit and generic algorithms for solving this kind of infinite games. A decision-making system in these complex economic settings require not only economic but also computational efficiency.

There are two research gaps to be bridged. First, as more and more auction mechanisms are introduced, it is useful to provide closed-form solutions. Second, urged by the need of industrial application, it is useful to design heuristic algorithms for large-scale problems which have not yet been solved analytically.

The flexible decision-making system should be designed to generate the optimal strategies for agents automatically, as illustrated in Figure 1.1.[4] To make use of the system, we need to specify our agent, the environment, and the market model.

- *Our agent:* Our agent is described by a utility function, a preference structure, and a distribu-

[4] This figure is a modified version of Figure 1, an architecture for trading agents, in [110].

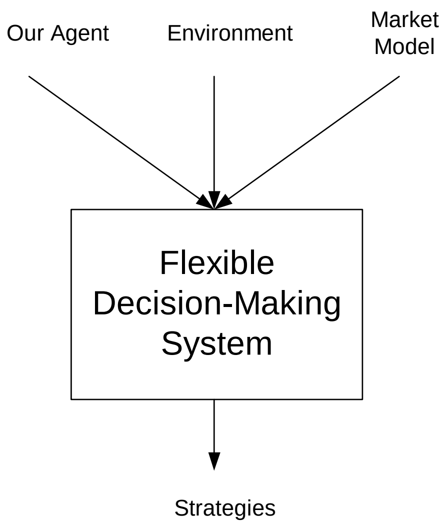

Figure 1.1: An architecture for a flexible decision-making system.

tion function, etc.

- *Environment:* The environment might be defined by a set of auction rules.
- *Market model:* We model other agents explicitly. The system knows about whether the agents know the strategies of each other and what strategies the other agents use. For example, some agents may use equilibrium strategies while the others use myopic strategies. A market model also includes whether the other agent's preferences and other information are known [110].

This thesis aims to provide answers to several issues. On the theory side, I provide closed-form solutions to FPSB auctions and sequential FPSB auctions. I also analyze the non-existence of equilibrium in two sequential auction models. On the algorithm side, I present a heuristic algorithm as part of a flexible decision-making system to compute solutions for sequential auctions with discrete bids [10].

In Chapter 2, I review some basic concepts of game theory and provide a brief survey on Nash equilibrium and its refinements. Chapter 3 provides a review of the state-of-art algorithms for computing equilibria.

In Chapter 4, I analyze the FPSB auctions, including sequential FPSB auctions, with discrete bids. An FPSB auction is a special case of sequential FPSB auctions when the number of items is equal to one. I discuss the existence and multiplicity of equilibria in the FPSB auctions with discrete bids under both complete information and incomplete information.

Chapter 5 focuses on a flexible decision-making system for sequential auctions with discrete bids. I present a heuristic approach using Monte-Carlo approximation. This system enables users to compute solutions for different sequential auction models more efficiently than existing algorithms.

In Chapter 6, I study the impact of information in sequential auctions when altering different information revelation policies. I prove the non-existence of pure-strategy symmetric equilibrium in both symmetric sequential first-price sealed-bid auctions and symmetric sequential Vickrey auctions.

Finally, Chapter 7 summarizes the contributions of the thesis and discusses directions for future work.

Chapter 2

Strategic Equilibria

In a broad sense, a game may refer to any social situation involving two or more individuals [78].[1] Individuals are also called players, agents, or decision-makers. Each individual is usually assumed to be *rational*, which implies that every player always maximizes his utility [28]. Game theory is the study of the noncooperation and cooperation between these rational players.

An *equilibrium* is defined as a state of a system that the system tends to move back to the same state when the system is perturbed from its original state. Equilibrium in a game is also called *strategic equilibrium*. Finding strategic equilibrium in games is a major task of game theorists.

2.1 Basic Concepts

For the sake of completeness, a brief review of the relevant definitions is provided.

An event E is *common knowledge* if all players know that E occurred, and all players know that all players know that E occurred, and so on, ad infinitum.

A game is of *certainty* if there is no stochastic events, which are typically characterized as a move by nature. If there is a move by nature, the game is said to have *uncertainty*.

A game is one of *symmetric information* if an agent's information state has the same elements as those of every other agent. Otherwise the game is said to be one of *asymmetric information*.

[1] There are different definitions and understandings of a game. In this thesis, we will use the definition by Myerson [78]. Thus, a game is a real world noncooperation or cooperation situation. This will also help us understand the definitions of game models. Under this definition, auctions are games.

A game is one of *incomplete information* if some or all of players lack full information about the timing of the game, the set of strategies, or the payoffs of players [28]. For example, nature moves first and is unobserved by at least one of the agents. Otherwise, the game is one of *complete information.*

A game is one of *perfect information* if each agent knows every action of the agents that moved before him at every point. Otherwise, it is one of *imperfect information*.

A strategic form game is *finite* if the number of players and the number of strategies is finite.

There are two kinds of games that have complete information but imperfect information. In the first scenario, the agents move simultaneously. In the second scenario, nature moves without revealing information immediately to all agents.

2.2 Game Models

A game model is a description of a game. Game models are also called game forms. Different forms abstract the game from a different perspective. To find a solution for a game, one builds a model of the game and then solves for an equilibrium of the model [23]. Due to variations in game models and copious equilibrium concepts, it is possible that we might have different answers as well as different specific solution procedures to the same game.

The two most important game forms are the *extensive* form and the *normal* (or *strategic*) form. In addition to these two forms, there are also the *agent normal* form and the *reduced normal* form. Due to limited space, we discuss only these four game models.[2]

2.2.1 Strategic Form and the Normal Representation

A strategic form has three elements: a set of *players*, A, a set of possible *(pure) strategies*,$\{S_i\}_{i \in A}$, and a set of *utility (payoff) functions*, $\{u_i\}_{i \in A}$. Thus, a strategic form game Γ can be denoted by $\Gamma = \{A, \{S_i\}_{i \in A}, \{u_i\}_{i \in A}\}$.

We let σ denote a *strategy profile* of the game. Let σ_i be a strategy profile of player i and $\sigma_{i,s}$ be the choice probability of each pure strategy, s, in player i's strategy set, S_i. If $\sigma_{i,s} \in \{0, 1\}$, we call σ a *pure strategy profile*. Otherwise, σ is a *mixed strategy profile*. For example, in a two-

[2]For more game models, interested readers may refer to [23, 78] or other game theory literature.

		Bidder 2 (v=2)	
		1	2
Bidder 1	1	0.75, 0.5	0,0
(v=2.5)	2	0.5, 0	0.25, 0

Table 2.1: A sealed bid auction game.

player game, each player has two actions, $\{a, b\}$. For a pure strategy profile in which Player 1 uses action b and Player 2 uses action a, the profile is written as $\sigma = \{b, a\}$. For a mixed strategy profile in which Player 1 has $1/2$ probability to use action a and Player 2 has $1/3$ probability to use action a, we have $\sigma = \{(1/2, 1/2), (1/3, 2/3)\}$.

As an example, a strategic form of a sealed bid auction, with two players and two bid values for each player, can be expressed as in Table 2.1, which is the *normal representation* of the game. In strategic forms, game theorists assume that players choose their strategies independently [80]. So, all the strategies in strategic form can be expressed in independent vectors. As a result, a strategic form game is equivalent to a normal form game.

2.2.2 Extensive Form

The extensive form is more richly structured than the normal form. Normally, we use a tree graph to depict an extensive form game. The tree consists of a set of *branches*, each of which connects two *nodes*. The first node is called *root*, which represents the beginning of the game and the bottom nodes are called *terminal nodes* and represent the end of the tree.

An extensive form game, Γ^e, includes six elements in which the first three elements are almost the same as in a strategic form game [23, 114].

1. A set of *players*, A, each of which has a player label.

2. The strategy space $S_i(\xi)$ of each player, i, at each information state (also called information set), ξ. An *information state* includes one player, the nodes at which the player has the same information, and the strategy space of the player at this information state. An information state includes two or more nodes if the player cannot distinguish between the situations represented by these nodes [78]. A node, d, may have several branches, which represent the feasible actions of the specific player. The feasible actions of the player can vary with position in the game tree.

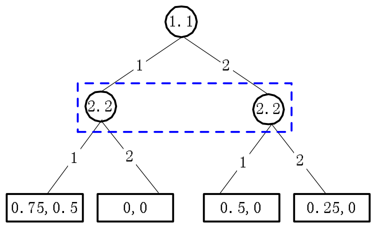

Figure 2.1: The extensive form of Table 2.1.

3. The players' payoff functions, $\{u_i\}_{i \in A}$. Normally, we label the payoff values at the terminal nodes.

4. The order of moves, i.e., who moves when.

5. The information state, ξ, of each player when she will move. We denote the set of information states as Ξ.

6. The stochastic events, which are encoded chance nodes annotated with their probabilities.

An example is shown in Figure 2.1. The first "1" in "1.1" identifies Player 1, where the second "1" identifies the first information state of Player 1. There are two "2.2"'s in the game which indicates that there is only one information state for Player 2 because Player 2 cannot observe Player 1's action. The single "1" and "2" along the branches are the feasible actions a player has in the information state.

A *behavioral strategy profile* in extensive forms refers to a probability over the set of possible strategies for each possible information state of each player. A behavioral strategy profile is very similar to a strategy profile defined in the normal form. The difference is that a behavioral strategy profile is related with information states. Let σ denote a behavioral strategy profile. Let $\sigma_{\xi,i}$ be a behavioral strategy profile of player i at information state ξ, and $\sigma_{\xi,i,s}$ be the choice probability of each pure strategy, s, in player i's strategy set, $S_i(\xi)$, at information state ξ. Thus, we have $\sigma = \{\sigma_{\xi,i}\}_{i \in A} = \{\sigma_{\xi,i,s}\}_{s \in S_i(\xi), i \in A}$. If $\sigma_{\xi,i,s}$ equals 0 or 1, we call σ a *pure behavioral strategy profile*. Otherwise, σ is a *mixed behavioral strategy profile*. For example, in Figure 2.2, $\sigma = \{2, b, d\}$ is a pure behavioral strategy profile, in which Player 1 selects action 1 at information state 1, Player 2 selects action b at information state 2, and Player 2 selects action d at information state 3. A mixed behavioral strategy looks like $\{2, (1/3, 2/3), (1/2, 1/2)\}$.

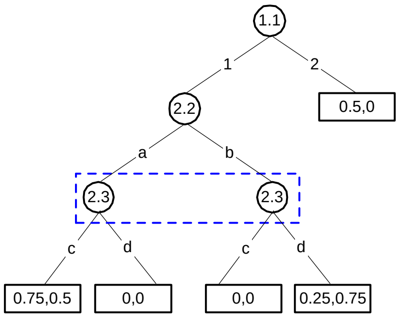

Figure 2.2: An extensive form game with imperfect recall.

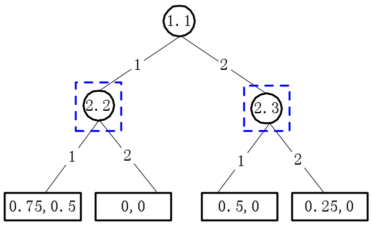

Figure 2.3: A perfect information game.

In extensive form games, *perfect recall* implies that a player will remember all the earlier information observed during the game, including her own past moves. Not all extensive form games have perfect recall. In the information state “2.3” of Figure 2.2, Player 2 cannot remember her past move and earlier information because she cannot recall which branch she just came from. Moreover, perfect recall is not equivalent to perfect information, in which each information state is a singleton. That is to say, perfect information is a stronger concept. For example, the extensive form game shown in Figure 2.1 is perfect recall because every player can remember the previous information and her past moves. However, Figure 2.1 is not a perfect information game because when Player 2 reaches state “2.2” she cannot determine which move player 1 made.

		Bidder 2			
		1,1	1,2	2,1	2,2
Bidder 1	1	0.75, 0.5	0.75, 0.5	0,0	0,0
	2	0.5, 0	0.25, 0	0.5, 0	0.25, 0

Table 2.2: The normal form representation of Figure 2.3.

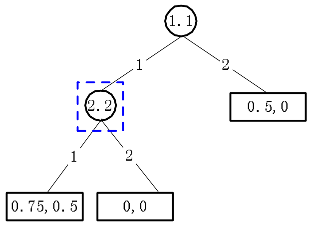

Figure 2.4: An extensive form game corresponding to the game in Table 2.3.

The Strategic Form Representation of Extensive Form Games

The normal form and extensive form are the two most common models for games. Even with these two models, we might have different results when we solve the same game.[3]

An extensive form can give us more information than a normal form game. It is possible to convert an extensive form game to strategic form, but we may lose information about the sequence of moves. To express this in normal form, we will assume that a player makes a complete contingent plan in advance [23]. Let those strategies in Γ^e be the pure strategies in Γ and let payoff functions be the same. Thus, the strategic form of the game in Figure 2.1 can be represented as in Table 2.1. In Figure 2.3, Player 2 has two information states corresponding to different moves by player 1. Totally, there are four pure strategy profiles, $\{1{,}1\}$,$\{1{,}2\}$,$\{2{,}1\}$,$\{2{,}2\}$, if represented in normal form. As a result, the normal form of the extensive form game in Figure 2.3 can be expressed as in Table 2.2.

It is not surprising that a normal form might have multiple extensive form representations. Consider Figure 2.4. The normal form representation of this game is shown in Table 2.3, which is very similar to Table 2.1. The only difference lies in that the payoff functions are the same when

[3]However, some researchers argued that different models of a game should provide the same solution. The reduced-normal form is a result of this discussion.

		Bidder 2	
		1	2
Bidder 1	1	0.75, 0.5	0,0
	2	0.5, 0	0.5, 0

Table 2.3: A sealed bid auction game corresponding to the game in Figure 2.4.

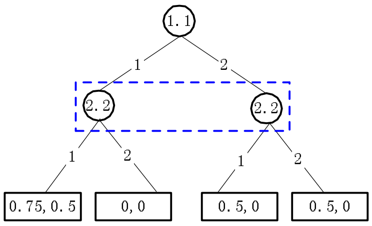

Figure 2.5: A different extensive form game of Table 2.3.

Player 1 chose "1" in both tables. Now, let us define an extensive form representing from the normal form game as shown in Table 2.3. From the same normal form game, we may have two different extensive form games, which are illustrated in Figure 2.4 and Figure 2.5. This interesting phenomenon confirms that we might lose some information when a normal form is represented from an extensive form.

Agent Normal Form

Defined by Selten [94], the *agent-normal form* representation is a modification of the extensive form, in which each information state in an extensive form game is associated with a different "temporary" agent. Those temporary agents share the same payoffs with the original agent. Some literature [23] also calls agent-normal form the *agent strategic form* or *multiagent representation form*. Similar to the relation of the normal form to the extensive form, a *multiagent representation* is a game in strategic form representing the corresponding "temporary agent" extensive form game.

To show the difference between the normal representation and the multiagent representation, consider the game in Figure 2.3. We have already shown the corresponding normal form in Table 2.2. The multiagent representation of Figure 2.3 is shown in Table 2.4. The set of players in

		Bidder 2			
		1		2	
		Bidder 3		Bidder 3	
		1	2	1	2
Bidder 1	1	0.75,0.75,0.5	0.75,0.75,0.5	0,0,0	0,0,0
	2	0.5,0.5,0	0.25,0.25,0	0.5,0.5,0	0.25,0.25,0

Table 2.4: The multiagent representation form of the game in Figure 2.3.

Table 2.2 is $A = \{1, 2\}$; the strategy profiles are $S_1 = \{1, 2\}$ and $S_2 = \{(1, 1), (1, 2), (2, 1), (2, 2)\}$; and the payoff functions are shown in Table 2.2. In comparison, the set of players in Table 2.4 is $A = \{1, 2, 3\}$; the strategy profiles are $S_1 = \{1, 2\}$, $S_2 = \{1, 2\}$, and $S_3 = \{1, 2\}$; and the payoff functions are shown in the Table 2.4. These two representations may result in different solutions as shown in a later discussion on perfect equilibrium, in the sense that the multiagent representation form is to rule out correlation between the "mistakes" of the same player in different stages of the game [94].

Reduced Normal Form

A reduced normal form, G, is a strategic form in which all pure strategies of a player that are convex combinations of other pure strategies of the same player have been deleted [43]. Let us examine some prerequisite concepts.

Definition 2.2.1. *Given any two strategies, c_i and d_i, in the strategy set S_i of player i, c_i and d_i are said to be* payoff equivalent *if and only if for all $s_{-i} \in S_{-i}$ and $j \in A$*

$$u_j(s_{-i}, c_i) = u_j(s_{-i}, d_i).$$

For example, the normal representation of the game in Figure 2.6 is shown as in Table 2.5. From the definition, we know that $2a$, $2b$, and $2c$ are payoff equivalent to each other. When two strategies are payoff equivalent, a player would be indifferent between them. In other words, we may replace the set of payoff equivalent strategies with a single strategy. A normal representation, in which we replace all sets of payoff equivalent strategies with a single strategy in the corresponding sets is called a *purely reduced normal representation*. The purely reduced normal representation of Figure 2.6 is shown in Table 2.6.

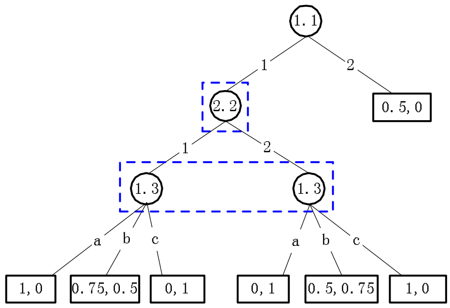

Figure 2.6: An extensive form game that can be reduced using reduced normal form.

		Bidder 2	
		1	2
	1a	1,0	0,1
	1b	0.75,0.5	0.5,0.75
Bidder 1	1c	0,1	1,0
	2a	0.5, 0	0.5, 0
	2b	0.5, 0	0.5, 0
	2c	0.5, 0	0.5, 0

Table 2.5: The normal form representation of Figure 2.6.

		Bidder 2	
		1	2
	1a	1,0	0,1
	1b	0.75,0.5	0.5,0.75
Bidder 1	1c	0,1	1,0
	2.	0.5, 0	0.5, 0

Table 2.6: The purely reduced normal form representation of Figure 2.6.

		Bidder 2	
		1	2
	1a	1,0	0,1
Bidder 1	1c	0,1	1,0
	2.	0.5, 0	0.5, 0

Table 2.7: The fully reduced normal form representation of Figure 2.6.

Definition 2.2.2. *A strategy e_i in S_i is randomly redundant if and only if it is a convex combination of the other pure strategies, where there is a probability distribution σ_i in $\Delta(S_i)$, the set of all randomized strategies for Player i, such that $\sigma_i(e_i) = 0$ and for all $s_{-i} \in S_{-i}$ and $j \in A$*

$$u_j(s_{-i}, e_i) = \sum_{d_i \in S_i} \sigma_i(d_i) u_j(s_{-i}, d_i).$$

For example, the strategy $1b$ is redundant because its payoff function can be expressed by a convex combination of strategies $1a$ and $1c$. A purely reduced normal representation is a fully reduced normal representation if it deletes all the randomly redundant strategies. The fully reduced normal representation of Figure 2.6 is shown in Table 2.7. Unless specified otherwise, we regard the reduced normal representation as the fully reduced normal representation.

2.3 Equilibrium Concepts

A *static (simultaneous) game* is one in which players will move simultaneously, without knowledge of the strategies that are being chosen by other players. A static game can be easily modeled as a normal form game.

A *dynamic game* will specify the order of moves. Unlike static games, players have at least some information about the choices made on past moves. The extensive form is usually used to express a dynamic game.

There are three milestones in the history of equilibrium concepts. The first one is the development of *Nash equilibrium* for static games with complete information. Nash equilibrium is the most famous and the most important equilibrium concept in game theory. John Nash was the first person to formally define the equilibrium of a non-cooperative general-sum game [80, 83].

		Player 2	
		Head	Tail
Player 1	Head	1,-1	-1,1
	Tail	-1,1	1,-1

Table 2.8: A game of matching pennies.

The second milestone was the introduction of *subgame perfect equilibrium* for dynamic games with complete information. Selten [94] refines the concept of Nash Equilibrium to subgame perfect equilibrium which can be applied to dynamic games and is computed using backward induction. The third milestone is Harsanyi's Bayesian equilibrium [31] which enables the agents in incomplete information games to choose strategies conditionally based on the perceptions of what the other agents are likely to do.

We discuss these three milestone equilibria together with other important equilibria in the following.

2.3.1 Nash Equilibrium

A *Nash equilibrium* describes a state of a multi-agent system in which no one can benefit by unilaterally changing her strategy. For example, in the game of Figure 2.1, the strategy profiles $\{1, 1\}$ and $\{2, 2\}$ are Nash equilibria.

We may explain the Nash equilibrium in another way. Suppose there is one agreement that all players promised to comply with prior to the game. This agreement is *self-enforcing* (or strategically stable) if no one would prefer to deviate and choose some strategy other than that specified in the agreement. Thus, to be self-enforcing, it is necessary that the agreement form a Nash equilibrium [49].

It is worth noting that, at the beginning of game theory, "more attention was focused on the cooperative analysis that von Neumann favored" [80]. Together with von Neumann's "cooperative" game theory, Nash equilibrium provides a "complete general methodology" to analyze all games [80]. In fact, a cooperative game can be reduced to a non-cooperative games in which "the steps of negotiation become moves" in the non-cooperative game [83]. Thus, Nash equilibrium applies to both cooperative games and non-cooperative games.

It was Nash's main contribution to show that *every finite game has a Nash equilibrium*

			Female	
			Badminton	Movie
Male	Badminton		2,1	0,0
	Movie		0,0	1,2

Table 2.9: Battle of the Sexes.

[83]. However, not every game has pure strategy Nash equilibrium; some games may have only mixed strategy equilibria. A classic example is "matching Pennies", shown in Table 2.8. In this game, two players simultaneously announce heads or tails. If the announcement matches, Player 1 wins; otherwise, Player 2 wins. There is no pure strategy Nash equilibrium in this game. The only stable situation is that both players play randomly between their two possible pure strategies with probabilities $(1/2, 1/2)$.

A game may have more than one Nash equilibrium. This problem often makes it hard to predict which Nash equilibrium will be played, and may complicate the computation of Nash equilibrium. We call this problem *multiplicity*. To solve this problem, game theorists try to provide some basis for claiming one equilibrium is better than another.

An allocation is *Pareto efficient* if no agent can be better off without making the others worse off. One criterion to find "better" equilibria is to look for Pareto efficient outcomes within the set of Nash equilibrium. If there are any overlaps, we call these overlapping Nash equilibria *Pareto dominant*. For example, in the game of Figure 2.1, the strategy profile $\{1, 1\}$ is Pareto dominant $\{2, 2\}$, and we can argue that it is a better equilibrium. However, this variant is often not conclusive.

Focal-point Effect The *Focal-point effect* argues that some of the given multiple Nash equilibria will be more plausible due to special properties they have. For example, suppose there are two Nash equilibria in one game, one pure strategy and another one mixed strategy. Some theorists argue that it is more preferable for the players to play the pure strategy [78]. Another focal-point effect example is Battle of the Sexes shown in Table 2.9. $\{Badminton, Badminton\}$ and $\{Movie, Movie\}$ are two Nash equilibria. However, suppose the female has some priority in the relationship; both players may be able to determine that $\{Movie, Movie\}$ will likely be the final result.

A *(strongly) dominant strategy* is the only optimal strategy for an agent no matter what strategies the other agents choose. If every agent has a dominant strategy, the set of these dom-

		prisoner 2	
		Testify	Conceal
prisoner 1	Testify	-3,-3	0,-5
	Conceal	-5,0	-1,-1

Table 2.10: A game of prisoner's dilemma.

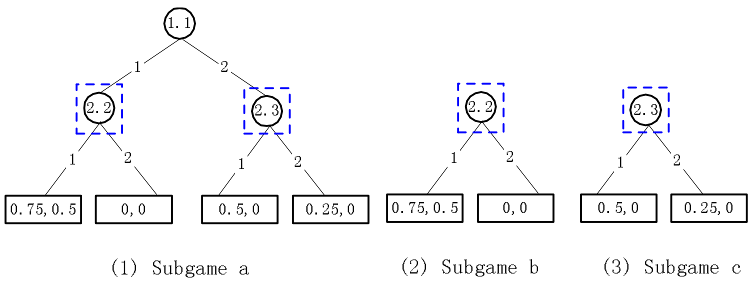

Figure 2.7: Subgame illustration of the game in Figure 2.3.

inant strategies is called a *dominant equilibrium*. A dominant equilibrium is a Nash equilibrium. A famous example is the Prisoner's Dilemma game. As shown in Table 2.10, if both prisoners do not testify, they each get -1 rewards; if both testify, they each get -3 rewards. If one testifies and the other does not, the former one gets 0 reward while the latter one gets -5. In the Prisoner's Dilemma, each player has a dominant strategy. However, the resulting equilibrium is Pareto dominated by an alternate outcome in which each player chooses the dominated strategy. It turns out that $\{testify, testify\}$ is the dominant strategy. It is worth noting that there is no Pareto dominance among equilibria in this game because there is only one Nash equilibrium. In fact, $\{testify, testify\}$ is the only solution which is not Pareto efficient.

A dominant strategy is a very strict condition of Nash strategy. In most cases, there are no, or only partially dominant strategies. After eliminating the dominated strategies, the remaining agents may find that their strategies become dominant in the reduced game. If this process can be continued until every agent eliminates all but one strategy, the game is *dominance solvable*. For example, in the game of Figure 2.1, the strategy profile $\{1, 1\}$ is a dominance solvable.

2.3.2 Subgame Perfect Equilibrium

A subgame is a component of a game. Let x be a node of an extensive form game, Γ^e. Let $g(x)$ be the set of all nodes and branches that follow x, including the node x itself. The node x is a *subroot* if, given any other node, $x^{'}$, in Γ^e, which happens at the same time at x or thereafter, either $g(x^{'}) \cap g(x) = \emptyset$ or $g(x^{'}) \subseteq g(x)$. We refer to $g(x)$ as a subgame, γ_x^e, of Γ^e. $\Gamma^e$is a subgame itself. For example, the game in Figure 2.3 has three subgame, as shown in Figure 2.7. In another example shown in Figure 2.1, the only subgame is the game itself. Because the information state of Player 2 cannot be separated, and the closest root of both nodes of this information state is the root of information state 1.

A behavioral strategy profile is a *subgame perfect equilibrium* if it introduces a Nash equilibrium to every subgame [94]. If there is more than one subgame in Γ^e, we may find an *equilibrium path*, in which, the restriction of the behavioral strategies to each subgame is an equilibrium. Let us look at the game in Figure 2.7 again. The Nash equilibrium of subgame b is $\{1\}$. In subgame c, either $\{1\}$ or $\{2\}$ could be the equilibrium solution for Player 2. Regardless of which one Player 2 picks, or whether she chooses a convex combination in subgame c, Player 1 will choose strategy $\{1\}$ in subgame a. Thus, there is only one subgame perfect equilibrium in this game, rather than two Nash equilibria as shown in Figure 2.3.

The concept of subgame perfect equilibrium is stronger than Nash equilibrium. If there is only one subgame in Γ^e, each subgame perfect equilibrium is a Nash equilibrium, and vice versa. However, if there is more than one subgame in Γ^e, the set of Nash equilibrium is a superset of subgame perfect equilibrium while every subgame perfect equilibrium is a Nash equilibrium.

Problems in Subgame Perfect Equilibrium Subgame perfect equilibrium is normally computed by *backward induction*, in which we solve the subgames at the leaves and work our way up the tree. The problem with backward induction is that it assumes the predictions of the agents' behaviors at the end of a game are credible even at the beginning [78]. An alternative method, introduced by Fudenberg, Kreps, and Levine [22], is *forward induction*. Forward induction requires an agent make decisions based on the information available in the earlier part of the game. For example, consider the game in Figure 2.8. There are two subgame perfect equilibria in this game, $\{2, b, d\}$ and $\{1, a, c\}$. However, the strategy profile $\{2, b, d\}$ will give Player 1 a lower payoff than the payoff gained by choosing 1 before the subgame. Thus, it is reasonable for Player 2 to infer that the only reason that Player 1 will choose move 1 at the first stage is that she expects $(0.75, 0.25)$

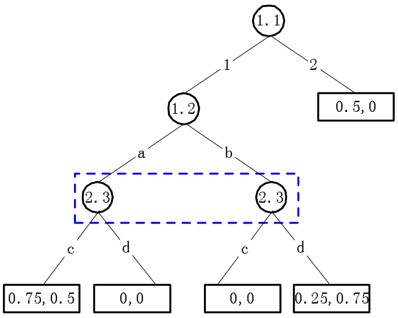

Figure 2.8: A game of forward induction.

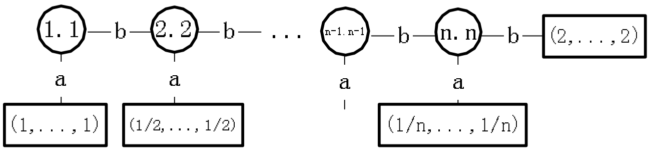

Figure 2.9: An n-player game illustrating subgame perfect equilibrium.

to happen in the subgame. As a result, Player 2 should choose c in the subgame, so $\{1, a, c\}$ is the only equilibrium in the forward induction.[4] Thus, forward induction may provide different solutions from backward induction. See page 192 in [78] for more examples.

Backward induction is not the only problem in subgame perfect equilibrium. Like Nash equilibrium, subgame perfect equilibrium assumes that all players are perfectly rational. Thus, all players expect an equilibrium in the whole game and the same equilibrium in every subgame. That is to say, subgame perfect equilibrium does not allow for imperfect play in the game. Consider the game in Figure 2.9. $\{b, b, ..., b\}$ is the only subgame perfect equilibrium and only forward induction equilibrium. However, there is one credible threat to $\{b, b, ..., b\}$. To show why, suppose that the first player has probability of $(1 - P)$ to choose a. If all players have the same probability to do so, the overall probability that all players will play b is P^n. If n is large, P^n will be small. So, in this sense, $\{b, b, ..., b\}$ will not be a good solution. As a re-examination of his subgame perfect equilibrium

[4] A similar discussion can also be applied to sequential equilibrium. However, there will be one more equilibrium in sequential equilibrium, $\{1, (0.25a + 0.75b), (0.6c + 0.4d)\}$, which is a mixed strategy profile. But, $\{1, a, c\}$ is still the only forward induction equilibrium.

		Bidder 2			
		Low Valuation		High Valuation	
		c	d	c	d
Bidder 1	a	0.75, 0.5	0,0	0.75, 1	0, 1
	b	0.5, 0	0.25, 0	0.5, 0	0.25, 0.5

Table 2.11: An incomplete information sealed bid auction.

concept, Selten [94] introduces a small "mistake" for every possible move of all players. We will touch this concept in sequential equilibrium and (trembling hand) perfect equilibrium.

2.3.3 Bayesian Equilibrium

Nash equilibrium assumes complete information. Difficulties arise in games of incomplete information in which players do not know each other's characteristics and hence the payment functions are no longer common knowledge. We illustrate this problem by discussing a two-person sealed bid auction, illustrated in Table 2.11. In this game, if bidder 2 has a low valuation of one item, $\{$low, low$\}$ is a weakly dominant Nash equilibrium. However, if bidder 2 has a high valuation, $\{$high, high$\}$ is a weakly dominant Nash equilibrium. So, whether bidder 1 chooses low or high will depend on whether bidder 2 has low valuation or high valuation.

Harsanyi [31] demonstrates that an incomplete information game can be transformed into a game with imperfect information. This kind of transformation is called *Harsanyi transformation*, in which an incomplete information game is replaced by a game where nature moves first (and chooses the players' types). As a result, we may have many complete information games with probabilities in accordance with the types of the players.

Let T_i be a set of possible types of player i, and T_{-i} denote all possible combinations of types for the players other than i. Let t_i be a typical type in T_i, and let t_{-i} be any possible combination of types for the players other than i. Let p_i be a probability function from T_i to $\Delta(T_{-i})$, which is the set of probability distributions over T_{-i}. We define $p_i(t_{-i}|t_i)$ for player i as the probability that the other players have t_{-i} while player i is in t_i. u_i denotes the utility function of player i. We define a *Bayesian game* as a profile,

$$\Gamma^b = \{A, \{S_i\}_{i\in A}, \{T_i\}_{i\in A}, \{p_i\}_{i\in A}, \{u_i\}_{i\in A}\}.$$

Thus, a strategy profile $\{\sigma_1, \sigma_2, \sigma_3, ...\}$ is a *Bayesian equilibrium* if the strategy of each player is a best response conditional on expectations of others' best responses, and no one wants to move unilaterally.[5] Let $\sigma_i^*(.|t_j)$ denote the best response strategy profile of player i given t_j. We have

$$\sigma_i^*(.|t_i) \in \arg\max_{s_i \in S_i} \sum_{t_{-i} \in T_{-i}} p_i(t_{-i}|t_i) u_i(s_i(t_i), s_{-i}(t_{-i})), (t_i, t_{-i})).$$

To see how to calculate a Bayesian equilibrium, let us consider the example in Table 2.11. In this game, player 1 has incomplete information while Player 2 has complete information. If Player 2 has a low valuation, Player 2 will play c because c is a weakly dominant strategy; otherwise, she will play d. We have

$$\sigma_1^*(.|t_2 = low) = a,$$
$$\sigma_1^*(.|t_2 = high) = b,$$
$$\sigma_2^*(.|t_2 = low) = c,$$
$$\sigma_2^*(.|t_2 = high) = d.$$

Suppose that the probability of Player 2's valuation being low is p. So, for Player 1, the expected payoff if she plays a is $0.75p$. If she plays b, the expected payoff is $0.5p+0.25(1-p) = 0.25p+0.25$. The critical value of p is $p = 0.5$. That is to say, if the probability of Player 2's valuation being low is less than 0.5, Player 1 will play b. For more examples, refer to page 215 in [23].

2.3.4 Perfect Bayesian Equilibrium

In Bayesian equilibrium, each player has a subjective probability distribution over the possible types of the other players. We refer to these subjective probability distributions as *prior beliefs*. The players do not modify the prior beliefs in the process of the game. However, in multi-stage games, players have the opportunity to observe the outcome of previous stages, and it is reasonable to think that players will modify their prior beliefs in accordance with the new information. The updated belief is called the *posterior belief*.

Perfect Bayesian equilibrium is an extension of subgame perfect equilibrium to incomplete information games. To formally define perfect Bayesian equilibrium, we let a_i^ξ be the action of player i at an information state ξ.[6] Let $\widetilde{p}_i(t_{-i}|a_{-i}^\xi)$ be the posterior probability of t_{-i} given that

[5] A Bayesian equilibrium is also called a *Bayes-Nash equilibrium*, or *expectation equilibrium*.

[6] The difference between action and strategy is trivial and many researchers use these two words interchangeably. To be precise, an action refers to a strategy that is used by one player and then is observed by the others.

player i observes the other players' moves leading to the information state ξ. A behavioral strategy profile is a *perfect Bayesian equilibrium* if at each information state ξ, we have

1. A player's strategy conditional on t_i is a best response to the other players' best response. For all $i \in A$ and $\xi \in \Xi$, we have

$$\sigma^*_{\xi,i}(.|t_i) \in \arg\max_{s_i \in S_i(\xi)} \sum_{t_{-i} \in T_{-i}} \widetilde{p}_i(t_{-i}|a_i^{\xi}) u_i(s_i, s_{-i}, t_i).$$

2. $\widetilde{p}_i(t_{-i}|a_{-i}^{\xi})$ is updated from a_i^{ξ} and s_{-i} using Bayes' rule whenever possible.

Thus, a perfect Bayesian equilibrium is a set of behavioral strategies and beliefs such that strategies are optimal given the beliefs at any stage of the game. The beliefs are updated from prior beliefs, equilibrium strategies, and observed actions using Bayes' rule.

2.3.5 Sequential Equilibrium

In perfect Bayesian equilibrium, there is no explicit definition of posterior probability when the observation probability is zero. As a result, there is no explicit definition of those strategies off the equilibrium path [114]. In this sense, perfect Bayesian equilibrium cannot guarantee an equilibrium solution for every subgame. Selten [94] introduces a concept referred to as "trembling hand perfection" to capture the notion that players may make errors with small probabilities. The trembling hand is a vivid description of "slight mistake" in which a player will do something wrong because she cannot hold her hand firmly. By introducing trembling, we enable the game to reach every information state.

This concept is applied in both *(trembling hand) perfect equilibrium* and *sequential equilibrium*. We introduce sequential equilibrium at first, because sequential equilibrium is simpler and normally easier to compute.

Kreps and Wilson [48] define (σ, μ) as an assessment, where σ is a behavioral strategy profile and μ is a set of beliefs at all information states. Let Σ be the set of all σs. $\sigma_{i(\xi)}$ denotes the strategy profile of player i at information state ξ and $\sigma_{-i(\xi)}$ denotes the strategy profile of all players except i at information state ξ. $u_{i(\xi)}$ denotes the utility of player i at ξ. Let $\mu_i(\xi)$ be the posterior probability distribution set of player i at the information state, ξ. We use Ψ to denote the set of all (σ, μ). So, a pair (σ, μ) is a *sequential equilibrium* if

1. (σ, μ) is sequential rational, that is, for every information state ξ,

$$u_{i(\xi)}(\sigma|\xi, \mu(\xi)) \geq u_{i(\xi)}((\sigma^{'}_{i(\xi)}, \sigma_{-i(\xi)})|\xi, \mu(\xi)),$$

   for all $i \in A$, $\xi \in \Xi$ and $\sigma^{'} \in \Sigma$.

2. (σ, μ) is consistent if there exists a sequence of strictly mixed (behavioral) strategy $(\hat{\sigma}^k)_{k=1}^{\infty}$, and associated beliefs $(\mu^k)_{k=1}^{\infty}$ determined by Bayes' rule, such that

$$(\sigma, \mu) = \lim_{k \to \infty} (\hat{\sigma}^k, \mu^k).$$

There are two points worth noting. First, players will adhere to the equilibrium profile σ at any information state including those off the equilibrium path. This is the same as subgame perfect equilibrium. Secondly, the behavioral strategies σ can be pure strategies, where (σ, μ) are limits of mixed strategies and associated beliefs. For example, consider a simple game in which Player 1 has a two-action strategy space. Suppose that $\sigma_1 = (1, 0)$ is the only pure strategy sequential equilibrium for Player 1. As required by the trembling hand property, we let $\hat{\sigma}_1^k = (1 - \epsilon_k, \epsilon_k)$. When $k \to \infty$, we have $\epsilon_k \to 0$ and $(\hat{\sigma}_1^k, \mu^k) \to (\sigma_1, \mu)$.

2.3.6 Perfect Equilibrium

First, let us discuss perfect equilibrium in strategic forms. We follow the definition on page 216 in [78]. Let $\Gamma = (A, (S_i)_{i \in A}, (u_i)_{i \in A})$ denote any finite game in strategic form. Let $\Delta(S_i)$ denote the set of all probability distributions on S_i and $\Delta^0(S_i)$ denote the set of all probability distributions on S_i that assign positive probability to every element in S_i. A strategy profile σ in $\times_{i \in A} \Delta(S_i)$ is a *perfect equilibrium* of Γ if and only if there exists a sequence $(\hat{\sigma}^k)_{k=1}^{\infty}$ such that

1. $\hat{\sigma}^k \in \times_{i \in A} \Delta^0(S_i)$,

2. $\sigma_i \in \arg\max_{s_i \in \Delta(S_i)} u_i(\hat{\sigma}^k_{-i}, s_i)$, and

3. $\lim_{k \to \infty} \hat{\sigma}_i^k(s_i) = \sigma_i(s_i)$, for all $i \in A$ and for all $s_i \in S_i$.

The first condition requires $\hat{\sigma}^k$ be a strictly mixed strategy profile in that every pure strategy of every player should have strictly positive probability. This is the same to the requirement in sequential equilibrium. The second condition asserts that σ_i is a best response strategy profile given every $\hat{\sigma}^k_{-i}$. This is stronger than in sequential equilibrium, which requires only that σ, the limit of

		Player 2 c	d
Player 1	a	0,1	0,1
	bx	-1,2	1, 0
	by	-1,2	2,3

Table 2.12: A game in strategic form.

$\hat{\sigma}^k$, is a Nash equilibrium. The third condition tells us that a perfect equilibrium is the converging limit of a sequence of Nash equilibria. This is a little bit different from sequential equilibrium, which puts more credit on posterior probability so that we may tell which beliefs are "plausible" [23].

However, for the purpose of perfectness, the strategic form is not an adequate representation of the extensive form [94]. In fact, a perfect equilibrium in strategic form may not even be a subgame perfect equilibrium due to "difficulties which may arise with respect to unreached parts of the game "[94]. To see why, let us look at the example in Table 2.12. In this game, $\{by, d\}$ is the only subgame perfect equilibrium. It is easy to understand that $\{by, d\}$ is also a perfect equilibrium. However, $\{a, c\}$ is also a perfect equilibrium. Suppose Player 1 will play a and, with some small probability, ϵ, tremble to bx and by, Player 2's expected payoff is $1 + 2\epsilon$ if Player 1 plays c and is $1 + \epsilon$ if she plays d. Since $1 + 2\epsilon > 1 + \epsilon$, $\{a, c\}$ is a perfect equilibrium. Some argue that the mistakes happened in different stages of subgame may be correlated [114]. To remove such "difficulties", Selten introduce "agent normal form as a more adequate representation of games with perfect recall" [94], which requires that an agent behaves independently in different stages such that an agent in a different stage looks like a different agent. Selten showed that every perfect equilibrium is always subgame perfect in agent normal form games, but the reverse may not hold.

2.3.7 Proper Equilibrium in Strategic Form

Proper equilibrium is a stronger equilibrium than perfect equilibrium, which requires a strictly positive probability for every pure strategy, but any "trembling" strategy is assigned an arbitrarily small probability. Proper equilibrium furthermore requires that any pure strategy that would be a mistake for a player is assigned a much smaller probability than any other strategy [78]. Myerson [77]

		Player 2	
		c	d
Player 1	ax	5,5	5,5
	ay	5,5	5,5
	bx	7,7	4,0
	by	0,0	3,3

Table 2.13: A normal form game.

defines that σ is an ϵ-proper equilibrium if and only if

1. $\sigma \in \times_{i\in A}\Delta(S_i)$,

2. For all $c_i, e_i \in S_i$,

$$\text{if } \mu_i(\sigma_{-i}, [c_i]) < \mu_i(\sigma_{-i}, [e_i]), \text{then } \sigma_i(c_i) \leq \epsilon\sigma_i(e_i).$$

A randomized-strategy profile σ in $\times_{i\in A}\Delta(S_i)$ is a proper equilibrium of Γ if and only if there exists a sequence $(\epsilon(k), \sigma^k)_{k=1}^{\infty}$ such that

1. For all k, σ^k is an ϵ-proper equilibrium,

2. $\lim\limits_{k\to\infty} \epsilon(k) = 0$, for all $k \in \{1, 2, 3, ...\}$, and

3. $\lim\limits_{k\to\infty} \sigma_i^k(s_i) = \sigma_i(s_i)$, for all $k \in \{1, 2, 3, ...\}$, for all $i \in A$ and for all $s_i \in S_i$.

As proved by Myerson [77], every proper equilibrium is a perfect equilibrium, but not vice versa. Consider the game in Table 2.13, the strategy profile

$$\{(1 - 7\epsilon)[ax] + \epsilon[ay] + \epsilon[bx] + 5\epsilon[by], (1 - \epsilon)[c] + \epsilon[d]\}$$

is an ϵ-perfect equilibrium. So, $\{ax, d\}$ is ϵ-perfect, as long as $0 < \epsilon < 1/3$. However, $\{ax, d\}$ is not an ϵ-proper equilibrium. The reason is that by is a worse mistake than bx for Player 1 because $0\epsilon + 3(1 - \epsilon) < 6\epsilon + 4(1 - \epsilon)$. The ϵ-properness condition requires that $\sigma_1(by)/\sigma_1(bx)$ must be no more than ϵ [78]. As a fact of matter, $\{bx, c\}$ is the unique proper equilibrium in this game as long as $\epsilon < 2/3$. This can be justified by the form

$$\{(1 - 2\epsilon - \epsilon^2)[bx] + \epsilon[ax] + \epsilon[ay] + \epsilon^2[by], (1 - \epsilon)[c] + \epsilon[d]\}.$$

		Player 2			
		c		d	
		Bidder 3		Player 3	
		e	f	e	f
Player 1	a	0,0,0	0,0,2	0,2,0	2,0,0
	b	0,2,0	2,0,2	2,0,2	0,2,0

Table 2.14: A 3-person game.

2.3.8 Persistent Equilibrium in Strategic Form

In the definitions of perfect equilibrium and proper equilibrium, trembles are forced when some of the pure strategies have zero probability. Thus, given that no pure strategy has zero probability, a Nash equilibrium is always perfect and proper [38]. However, this kind of strategy combination, referred to as an *inner combination*, is not always immune against trembles, and thus could be unstable. Recall the “battle of the sexes” game in Table 2.9. $\{Badminton, Badminton\}$, $\{Movie, Movie\}$, and $\{(1/2Badminton + 1/2Movie), (1/2Badminton + 1/2Movie)\}$ are the only three Nash equilibria, which are also perfect and proper. However, $\{(1/2Badminton + 1/2Movie), (1/2Badminton + 1/2Movie)\}$ does not have neighborhood stability since any trembles, like $\{((1/2+\epsilon)Badminton + (1/2-\epsilon)Movie), (1/2+\epsilon)Badminton + (1/2-\epsilon)Movie)\}$ will cause it to shift to $\{Badminton, Badminton\}$.

Here are some prerequisite definitions. A *retract* of the game Γ is defined as a subset R of S if $R = \times_{i \in A}\Delta(R_i)$, with each R_i being a non-empty closed convex subset of $\delta(S_i)$. A retract R is absorbing $\hat{S}$ if $BR(\sigma) \cap R \neq 0$ given that $\hat{S}$ is a set of mixed strategies $\hat{S} \subseteq S$ and every $\sigma \in \hat{S}$. That is, for every player i there is a $\tau_i \in R_i$ such that τ_i is a best response of player i to $\sigma \in R$ [38].

A retract R is *persistent* if it is a minimal absorbing retract. A strategy profile σ is a persistent equilibrium if σ is a Nash equilibrium and is a persistent retract [38].

Kalai and Samet prove that any finite game in strategic form has a persistent equilibrium which is perfect and proper. However, there may exist some persistent strategies which are not proper [38]. Let us look at an example, as shown in Table 2.14. In this game, every strategy profile is persistent. However, the strategy $\{a, c, e\}$ is not perfect.[7]

[7] See page 44 in [38] for the proof.

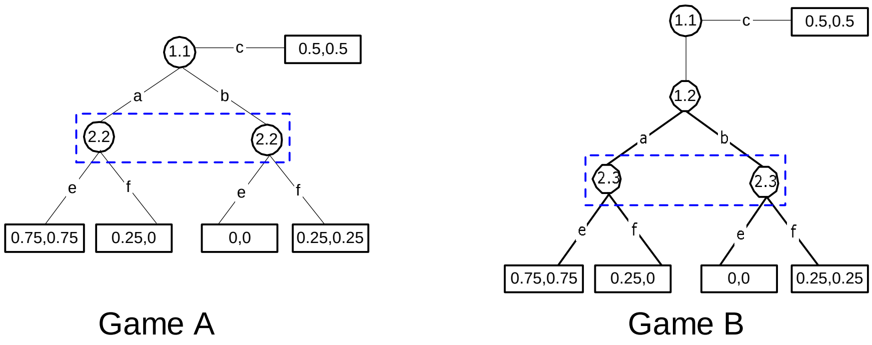

Figure 2.10: Two extensive games with the same reduced normal form.

2.3.9 Stable Equilibrium

As we know, a game in normal form could have different equilibrium solutions, when compared to those in extensive form. Consider the two games in Figure 2.10.[8] The reduced normal form game of these two games are the same; however, $\{c, f\}$ and $\{a, e\}$ are the perfect equilibria of Game A, while $\{a, e\}$ is the unique perfect equilibrium of game B. At the same time, backwards induction of the extensive form and the iterated dominance of the normal form do not give us the same "strategically stable equilibrium" [43]. Consider Game B in Figure 2.10 again. Strategy b of player 1 is strongly dominated by strategy c. So, $\{a, e\}$ should be the only equilibrium in iterated elimination of dominated strategies.

Stable equilibrium is a concept developed by Kohlberg and Mertens [43] to solve the above discrepancies. A reduced normal form, G, is where all pure strategies that are convex combinations of other pure strategies have been deleted [43]. Kohlberg and Mertens point out that a strategically stable equilibrium should depend only on the reduced normal form of the game. A strategically stable set of equilibria of G must contain a strategically stable set of equilibria of any $G^{'}$, which is obtained from G by a deletion of any dominated strategy [43].

We define S as a closed set of Nash equilibrium of G, if for any $\epsilon > 0$ there exists some $0 < \delta_0 \leq 1$, such that the perturbed game, where every strategy s of player i is replaced by $(1-\delta_i)s+\delta_i\sigma_i$, has an equilibrium ϵ-close to S, for any completely mixed strategy vector $\sigma_1, ..., \sigma_n$ (n players) and for any $\delta_1, ..., \delta_n$, ($0 < \delta_i < \delta_0$). A set of equilibria is stable in a game G if it is the minimal set of S [43].

[8]Game A and Game B in Figure 2.10 are similar to the games in Figure 2 and Figure 3 of [43].

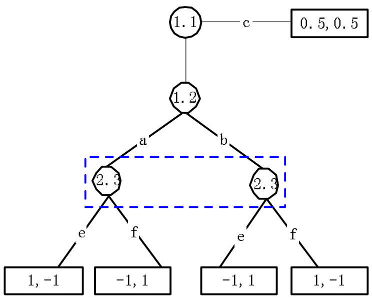

Figure 2.11: An extensive game corresponding to the game in Table 2.15.

		Player 2	
		e	f
	a	1,-1	-1,1
Player 1	b	-1,1	1,-1
	c	0.5	0.5

Table 2.15: A normal form game corresponding to the game in Figure 2.11.

Kohlberg and Mertens prove that, in iterated dominance, a stable equilibrium contains a stable set of any game obtained by eliminating dominated strategies. In forward induction, a stable equilibrium contains a stable set of any game obtained by a deletion of any strategy which is an inferior response to the equilibria of the set [43]. However, stable sets might not satisfy the backwards induction requirement. Consider the game in Table 2.15. There are two stable equilibria, $\{c, (1/4, 3/4)\}$ and $\{c, (3/4, 1/4)\}$. However, in the corresponding extensive form game in Figure 2.11, the only sequential equilibrium is $\{c, (1/2, 1/2)\}$. The cause of this problem might be because stable equilibrium uses a different game form other than normal form or extensive form representations. Similarly, it may not be a subset of proper or perfect equilibrium.

2.4 Summary

Using static and dynamic as one axis, and information as another, we classify these well-known equilibrium concepts in Table 2.16.

We learn that the results of computing different equilibria also depend on the game mod-

	Complete Information	Incomplete Information
Static	Nash Equilibrium	Bayesian Equilibrium
Dynamic	Subgame Perfect Equilibrium	Perfect Bayesian Equilibrium; Sequential Equilibrium; Perfect Equilibrium; Persistent Equilibrium; Stable Equilibrium;

Table 2.16: Categories of equilibria.

els. Usually, Nash equilibrium, proper equilibrium, and persistent equilibrium are solved in normal form. Subgame perfect equilibrium and sequential equilibrium are applied to extensive form games. Perfect equilibrium is discussed in agent normal form. And stable equilibrium is solved in reduced normal form. The choice of a game model to a specific application should depend on the needs of the scenario.

To conclude, we may depict a rough relationship picture among some of these equilibria, as shown in Figure 2.12. Normally, a stricter equilibrium concept is a subset of another equilibrium concept. The reason that persistent and stable equilibrium are not included in the picture lies in that they are not strictly a subset of proper or perfect equilibrium and could overlap with other concepts.

Simply, there is not an optimal equilibrium concept. Selection of an equilibrium concept to a specific game will depend on the properties of the game and the needs of the modeler. These refinements of equilibria provide many options, while, at the same time, they introduce different computational complexity to solving a game.

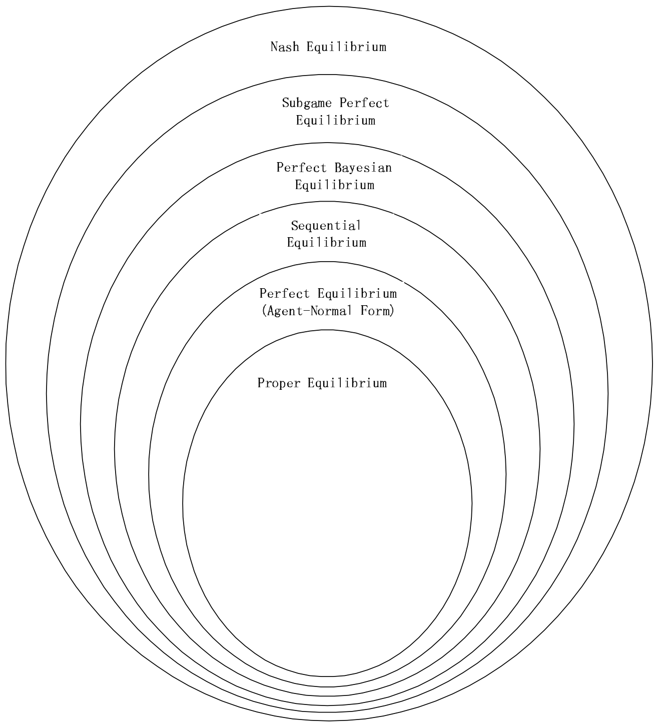

Figure 2.12: Relationship among different equilibria.

Chapter 3

Computing Equilibria

The concepts of Nash equilibrium and its refinements have been widely applied in economics, business, and other realms. Naturally, the computation of equilibria has drawn much attention. In general, the computational complexity of solving games is exponential. There are many papers focus how to solve 2-person games, and more recently there are more and more algorithms aiming to computing n-person games. To date, the solvable size of games has remained small. However, these algorithms are significant because many large size games can be approximated by smaller ones.

3.1 The Mathematics of Computing Nash Equilibrium

3.1.1 Nash Equilibrium as a Fixed Point of a Function

Nash used the fixed point theorem to prove the existence of equilibrium for finite, n-person games. Many algorithms for solving n-person games follow this idea and first find the fixed points. Scarf's algorithm was the first algorithm developed for approximating a fixed point by using algebraic sets [92]. This work was followed by many simplicial subdivision algorithms [99, 101, 113]. Another approach to approximate fixed point is using simplicial homotopy methods [14, 113]. However, the computational complexity of these algorithms, in the worst case, is exponential in the dimension and the number of digits of accuracy [33].

3.1.2 Nash Equilibrium as a Solution to Linear Complementary Problem

The Lemke-Howson algorithm was the first linear complementary problem algorithm to solve general sum, 2-person games [56, 66]. Modified versions of Lemke-Howson algorithm can be used to solve n-person games [88]. However, these algorithms need a non-linear component to deal with the transformation from the original form to the linear complementary problem. The Lemke-Howson algorithm has an exponential lower bound [76], and adding a non-linear transformation makes it even more computationally demanding.

Constant sum games are a special case of the class of 2-person games and are easier to solve. These games can be represented by primal-dual linear programs, which can be solved in polynomial time.

3.1.3 Other Mathematical Approaches

Here is a list of other mathematical approaches that can be used to compute equilibria. First, Nash equilibrium can be approximated as a solution to non-linear complementary problem. In comparison to linear complementary problem for 2-person games, non-linear complementary problem can be used for n-person games [66]. Second, Nash equilibrium is solved as a stationary point problem. The Kakutani fixed point theorem is implied by the stationary point theorem [113]. Third, Nash equilibrium is mapped to a semi-algebraic set. Fourth, Nash equilibrium is formulated as a minimum of a function on a polytope [66].

3.2 Computing a Sample Nash Equilibrium in Two-Person Games

3.2.1 Zero Sum Normal Form Games

Zero-sum normal form 2-person games are the simplest games in terms of computational complexity. The minimax algorithm is usually used to solve this class of games.

The MiniMax Algorithm

In a zero-sum, normal form game, payoff functions are common knowledge to both players. When a player wants to maximize her payoff, it is equivalent for this player to minimize the other player's payoff. In this sense, a minimax strategy is that both players want to minimize their maximum possible loss [84]. It has been proved that a pair of strategies is a Nash equilibrium if and only if it is minimax [66, 84]. We can use a linear program to describe the minimax problem for both players, and the solution to the primal-dual problem of the linear program is a Nash equilibrium.

To demonstrate how to solve the minimax and the primal-dual problem, let U_i be the payoff matrix of player i. Since the game is zero sum, $U_1 = -U_2$. To simplify, let $U = U_i$. Consider a mixed strategy, where P_i is the probability density vector among all pure strategies of player i, and $\sum_i P_i = 1$. Let s and t denote the additional scalar variables. As we will see, the primal problem is for player 1, while the dual problem is for the other player. The primal problem can be expressed as follows [84].

$$\begin{aligned}
&\phi = \min_{P_1,s} s \\
&\text{where} \\
&S = \{(P_1, s)|UP_1 \leq s1^n, 1^m P_1 = 1 \text{ and } P_1 \geq 0\}, \\
&S^* = \{(P_1, s) \in S|\phi = s\},
\end{aligned}$$

and, the dual problem is:

$$\begin{aligned}
&\psi = \min_{P_2,t} t \\
&\text{where} \\
&T = \{(-P_2, s)|P_2 U \leq t1^m, 1^n P_2 = 1 \text{ and } P_2 \geq 0\}, \\
&T^* = \{(P_2, s) \in S|\psi = t\}.
\end{aligned}$$

Furthermore, let $x = P_1/s$ and $y = -P_2/t$, the original primal-dual problem can be reduced to the primal problem:

$$\begin{aligned}
&\phi = \min_{x \in S_G} -1^m x \\
&\text{where} \\
&S_G = \{x \geq 0|Ux \leq 1^n\}, \\
&S_G^* = \{x \in S_G \in S|\phi_G = -1^m x\}.
\end{aligned} \tag{3.1}$$

and, the dual problem:

$$\begin{aligned}
&\psi = \min_{y \in S_G} -1^n y \\
&\text{where} \\
&T_G = \{x \geq 0| -yU \leq -1^m\}, \\
&T_G^* = \{y \in T_G \in S|\psi_G = 1^n y\}.
\end{aligned} \tag{3.2}$$

We define

$$E(P_1, P_2) \equiv P_2 U P_1.$$

A solution $\{P_1^*, P_2^*\}$ is a pair of strategies P_1^* and P_2^*, such that

$$E(P_1^*, P_2) \leq E(P_1^*, P_2^*) \leq E(P_1, P_2^*).$$

These minimax solutions are saddle points of the function E [84]. The solution $\{P_1^*, P_2^*\}$ is a Nash equilibrium.

Complexity Results

We can use the dual simplex algorithm to solve this primal-dual problem [84]. The complexity of dual simplex algorithm is similar to that of the simplex algorithm. The difference lies in that the dual simplex algorithm uses different criteria to pick the pivoting elements. Current LP-solvers can solve the linear program problem quickly, especially for sparse matrix LPs; however, in the worst case, the simplex method requires exponential time [40]. In 1984, Karmarkar developed the interior point method, which can solve the LPs in polynomial time. In summary, the minimax algorithm for the constant-sum, 2-person, normal form game is in the polynomial class.

3.2.2 General Sum Normal Form Games

The first and most well-known algorithm for solving the general sum, two-person normal form game is the Lemke-Howson algorithm [56], which reduces a normal form game to a linear complementarity problem.

Linear Complementarity Problem

A linear complementarity problem (LCP) consists of a set of inequalities and equations. The aim of the *linear complementarity problem* is to find a vector $z \in R^n$, such that

$$\begin{aligned} & z \geq 0 \\ & q + Mz \geq 0 \\ & z^T(q + Mz) = 0 \end{aligned} \tag{3.3}$$

given a vector $q \in R^n$ and a matrix $M \in R^{n \times n}$ [13]. There are many algorithms for solving linear complementarity problems.

The Lemke-Howson Algorithm

From the minimax algorithm, we know that Nash equilibria in a two-person game are equivalent to the minimax solutions. However, the minimax property no longer holds in the general sum two-person game because $U_1 \neq -U_2$. A pair of strategies (P_1^*, P_2^*) is a Nash equilibrium if and only if $E(P_1^*, P_2) \leq E(P_1^*, P_2^*)$ and $E(P_1, P_2^*) \leq E(P_1^*, P_2^*)$. These two equations are equivalent to

$$\begin{aligned} &P_1^T U_1 P_2^* \leq (P_1^*)^T U_1 P_2^*, \ \ and \\ &(P_1^*)^T U_2 P_2 \leq (P_1^*)^T U_2 P_2^*. \end{aligned}$$

We cannot reduce these equations to a primal-dual problem because they do not share the same utility function. However, we may construct two special functions so that we may convert them to a linear complementarity problem [13]. Without loss of generality, we assume that both U_1 and U_2 are positive matrices. If not, we can add a large scalar to make them positive. This alteration will not change the solution of the Nash equilibrium. We construct a linear complementarity problem as follows.

$$\begin{aligned} &u = -e_m + U_1 P_2 \geq 0, P_1 \geq 0, P_1^T u = 0, \\ &v = -e_n + P_1^T U_2 \geq 0, P_2 \geq 0, P_2^T v = 0, \end{aligned} \tag{3.4}$$

where e_m and e_n are two vectors whose components are all ones [13]. Let

$$q = \begin{bmatrix} -e_m \\ -e_n \end{bmatrix}, M = \begin{bmatrix} 0 & P_1 \\ P_2 & 0 \end{bmatrix}, \text{ and } z = \begin{bmatrix} u \\ v \end{bmatrix}. \tag{3.5}$$

3.4 and 3.5 are equivalent to 3.3. Suppose $(P_1^{'}, P_2^{'})$ is a solution to 3.4, the relation between $(P_1^{'}, P_2^{'})$ and the Nash equilibrium (P_1^*, P_2^*) is given by

$$P_1^* = P_1^{'}/(e_m)^T P_1^{'} \text{ and } P_2^* = P_2^{'}/(e_n)^T P_2^{'}. \tag{3.6}$$

Lemke and Howson prove that every solution of a non-degenerate LCP is a solution of the bimatrix game. The pivoting procedure and lexicographic degeneracy resolution are similar to those in linear programming.[1]

Complexity Results

The Linear complementarity problem is NP-hard [12] and thus there is no polynomial algorithm for the Lemke-Howson algorithm. Several algorithms for LCP are presented in [13], but all of them are computationally costly.

[1] A detailed description of the Lemke-Howson algorithm is on page 285 of [13]. A more general discussion of LCP, including degenerate cases, can be found in both [13] and [66].

3.3 Computing a Sample Nash Equilibrium in n-person Games

The algorithms for computing 2-person games cannot be directly extended to solving n-person games. A practical approach is to transform a game into a fixed point problem. The simplicial subdivision algorithms, including Scarf's algorithm, are used to find fixed points of the continuous function. Before we discuss the algorithms in detail, let us introduce some basic concepts in topology.

Definition 3.3.1. *Let $x^1, ..., x^m$ be vectors in $\mathbb{R}^n$ given that $\mathbb{R}^n$ is an n-dimensional Euclidean space. We say that $x = \sum_{i=1}^{m} \lambda_i x^i$ is a* linear combination *of $x^1, ..., x^m$ if $\lambda_i \in \mathbb{R}$ and $i \in I_m$, where I_m is an m element natural number finite set. $x = \sum_{i=1}^{m} \lambda_i x^i$ is an* affine combination *if x is a linear combination and $\sum_{i=1}^{m} \lambda_i = 1$. $x = \sum_{i=1}^{m} \lambda_i x^i$ is a* convex combination *if x is an affine combination and $\lambda_i > 0$, for all i. We say that $x^1, ..., x^m$ are* affinely independent *if the unique solution to $\sum_{i=1}^{m} \lambda_i x^i = 0$ is $\lambda_i = 0$, for all $i \in I_m$.*

Definition 3.3.2. *If $x^1, ..., x^{m+1}$ are affinely independent, the convex hull of $x^1, ..., x^{m+1}$ is an* m-dimensional simplex *or* m-simplex. *A k-simplex $\tau(< m)$ is called a* k-face, f, *or* k-dimensional face *of an m-simplex σ if all vertices of τ are vertices of σ. If a face f consists of a single point, it is called a* vertex. *If f is a half line or a line segment, it is called an* edge. *If f has a dimension one lower than the dimension of σ, it is called a* facet *of σ.*

Definition 3.3.3. *Let S denote a m-dimensional convex set in $\mathbb{R}^n$. A collection $\mathbb{T}$ of m-simplices is said to be a* triangulation *or* simplicial subdivision *of S if*

1. *$S = \bigcup_i \sigma_i$, for all $\sigma_i \in \mathbb{T}$*

2. *$\sigma_1 \cap \sigma_2 = \emptyset$ or $\sigma_1 \cap \sigma_2 = f$. There exists f, such that f is a common face of both σ_1 and σ_2, for all $\sigma_1, \sigma_2 \in \mathbb{T}$.*

3. *Let δ be a neighborhood of x, for all $x \in S$. There are only a finite number of simplices, belonging to $\mathbb{T}$, in δ.*

Definition 3.3.4. *Let $diam(\sigma)$ be the* diameter *of a simplex $\sigma \in \mathbb{T}$, where*

$$diam(\sigma) = \max\{\parallel x - y \parallel |x, y \in \sigma\}.$$

Let $mesh(\mathbb{T})$ denote the mesh size *of $\mathbb{T}$, where*

$$mesh(\mathbb{T}) = \sup\{diam(\sigma)|\sigma \in \mathbb{T}\}.$$

Definition 3.3.5. *A labeling rule $l : \mathbb{T}^0 \mapsto I_m$ is called a* proper labeling rule *if $x_i = 0 \Rightarrow l(x) \neq i$, for all $x \in \mathbb{T}^0$.*

A proper labeling is also called *Sperner proper labeling*. A possible proper labeling rule can be given by

$$l(y) = \min\{i \in I_n | f_i(x) \leq x_i > 0, f_i(x) \geq x_{i+1}\} \tag{3.7}$$

where $l(x_i) + 1 = 1$ if $l(x_i) = n$.

Definition 3.3.6. *Given a labeling rule $l : \mathbb{T}^0 \mapsto I_m$, an $(m-1)$-simplex σ with vertices $x^1, ..., x^m$ is a* completely labelled *simplex if all its vertices are differently labelled such that $\{l(x^i)|i \in I_m\} = I_m$. An $(m-1)$ or $(m-2)$-simplex σ is* almost completely labelled *if its vertices have at least all labels in I_{m-1}.*

Definition 3.3.7. *Two ordered set are said to be* adjacent *if they differ by at most one element. Thus, an almost completely labelled simplex σ is adjacent to at most two other simplices. A simplex is called a* terminal simplex *if there are at most one adjacent simplex. Let P be the set of simplices can be reached by σ. P is called a* loop *if there is no terminal simplex, or a* string *if there are two terminal simplices, or a* point *if there is only one terminal simplex.*

Theorem 3.3.1 (Brouwer Theorem). *Let D be a nonempty, compact, and convex subset of $\mathbb{R}$, and let $f : D \mapsto D$ be a continuous function. Then there exists at least one point x^* in D such that $f(x*) = x^*$ [113].*

Kakutani [1941] proves that a looser requirement can also result in a fixed point.

Theorem 3.3.2 (Kakutani Theorem). *Let D be a nonempty, compact, and convex subset of $\mathbb{R}$, and let $f : D \mapsto D$ be a upper semi-continuous mapping. Then there exists at least one point x^* in D such that $f(x*) = x^*$. A point-to-set mapping f is upper semi-continuous at the point x^* in D if for any convergent sequence $\{x^k | k \in \mathbb{N}\}$ of points in D, and for any convergent sequence $\{y^k | k \in \mathbb{N}\}$ with $y^k \in f(x^k)$ and limits y^*, it holds that $y^* \in f(x^*)$ [113].*

Nash proved the existence of Nash equilibrium in general games using the Kakutani Theorem.

Theorem 3.3.3 (Sperner Lemma). *Let D^i be a closed subset and let S^n denote a collection of closed subsets where $S^n = \bigcup\limits_{i=1}^{n} D^i$. Let $\mathbb{T}$ be a triangulation of S^n. Let $l : \mathbb{T}^0 \mapsto I_n$ be a proper labeling function for every vertex of $\mathbb{T}^0$. Then there exists at least one completely labelled simplex in $\mathbb{T}$ [113].*

The Sperner lemma guarantees a completely labelled simplex inside a triangulation using the Sperner proper labeling. In detail, the Sperner proper labeling rule $l : \mathbb{T}^0 \mapsto I_n$ in a triangle's triangulation has the following requirements [6]:

1. The vertices of the original triangle are labeled with three different labels, which are an element of $I_n, n = 3$.

2. Vertices of the triangulation that lie on a side, e.g. 1,2, should be labelled either 1 or 2.

3. There is no restriction on labeling of the vertices of any subsimplicies.

Using a proper labeling rule, e.g. equation 3.7, we may construct a function such that $\sum\limits_{i=1}^{n} f_i(x) = \sum\limits_{i=1}^{n} x_i = 1$. Since the Sperner lemma guarantees a complete labelled simplex inside the triangulation, we have $f_i(x^*) = x^*, i = 1, ..., n$. Hence x^* is a fixed point of f.

3.3.1 Scarf's Algorithm

The Brouwer theorem and the Sperner lemma prove the existence of at least one fixed point; however, they did not provide an approach to find the fixed points. Scarf was the first person to develop algorithms for this purpose. Scarf presented two algorithms for finding fixed points. One is in the framework of simplices; while the other one is organized around primitive sets [92, 93]. Labelled primitive sets in Scarf's algorithm are an analogue of labelled simplices. Although Scarf used primitive sets to explain the algorithm, simplicial division has become a more popular approach to solve the fixed point problem. We will mainly discuss the simplicial division algorithm here.

Basically, Scarf's algorithm starts with a randomly picked almost completely labeled $(n-2)$ simplex and then generates a finite sequence of adjacent almost completely labeled $(n-2)$ simplex. It stops at a completely labeled $(n-1)$ simplex.

The original Scarf's simplicial division algorithm is based on $K_2(m)$-triangulation, e.g. as shown in Figure 3.1.

Definition 3.3.8. *A triangulation is called* $K_2(m)$-triangulation *of* S^n *with grid size* m^{-1} *if it is a collection of all (n-1)-simplices* $\sigma(x^i, \pi)$ *with vertices* $x^1, ..., x^n$ *in* S^n *such that*

1. *each component has equal size.*

2. $\pi = (\pi(1), ..., \pi(n-1))$ *is a permutation of the elements in* I_{n-1}.

3. $x^{i+1} = x^i + m^{-1}t(\pi(i)), i \in I_{n-1}$*, where* $t(\pi(i)) = e(j+1) - e(j), j \in I_{n-1}$.

The triangulation for Scarf's algorithm, as shown in Figure 3.2 is a modification of $K_2(m)$-triangulation. The $K_2(m)$ is completely contained in the original triangulation. In this example, $n = 3$ and $m = 4$. We set the label of the original triangulation as $(i+1)$ mod n. The similar construction of subsimplices can be applied to the destination simplex if we want a finer approximation. The Scarf algorithm works as follows [113]:

1. Find a start simplex σ^0 such that σ^0 is a unique $(n-1)$-simplex in the original triangulation and has τ^0 as its facet, which is a $(n-2)$ simplex on the border. Let x^+ be the vertex of σ^0 that is not a vertex of of τ^0. Let $k = 0$.

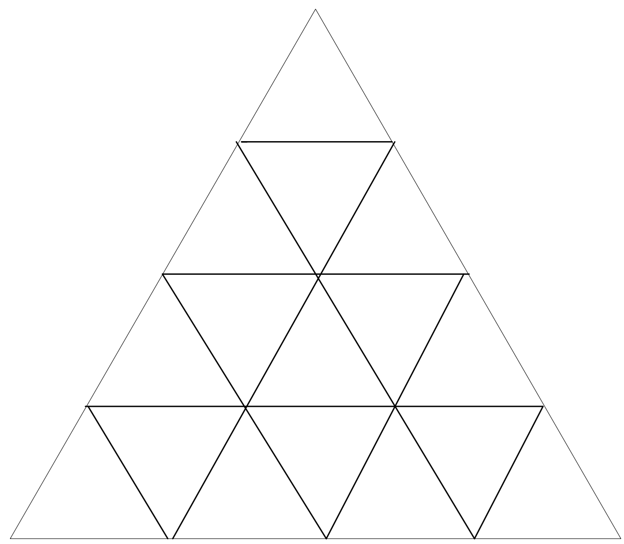

Figure 3.1: Illustration of K_2(m)-triangulation.

2. If $l(x^+) = n$, a completely labelled simplex σ^k has been found and the algorithm stops. Otherwise, $l(x^+)$ is the label of one other vertex of σ^k, say x^-. Let τ^{k+1} be the facet of σ^k opposite x^-.

3. Find the simplex σ^{k+1} adjacent to σ^k sharing τ^{k+1}. Let x^+ be the vertex of σ^{k+1} that is not a vertex of of τ^{k+1}. Let $k = k + 1$ and go to (2).

3.3.2 Transformation from a Game to a Fixed Point Problem

Scarf's algorithm did not provide the link from a game to a fixed point problem. To enable that, we need to define a transformation mapping the fixed points to the Nash equilibria.

Let s_{ij} be the j^{th} pure strategy of player i. Let p_{ij} be the probability over s_{ij}, p_i be the probability function of player i, p_{-i} be the probability of all players except player i, and p be the overall probability function. Let $\Delta = \{p | \sum_i p_i = 1, p_i > 0\}$ and let $f : \Delta \mapsto \Delta$. We define

$$g_{ij}(p) = \max[u_i(s_{ij}, p_{-i}) - u_i(p), 0] \text{ and } f_{ij}(p) = \frac{p_{ij} + g_{ij}(p)}{1 + \sum_j g_{ij}(p)}. \tag{3.8}$$

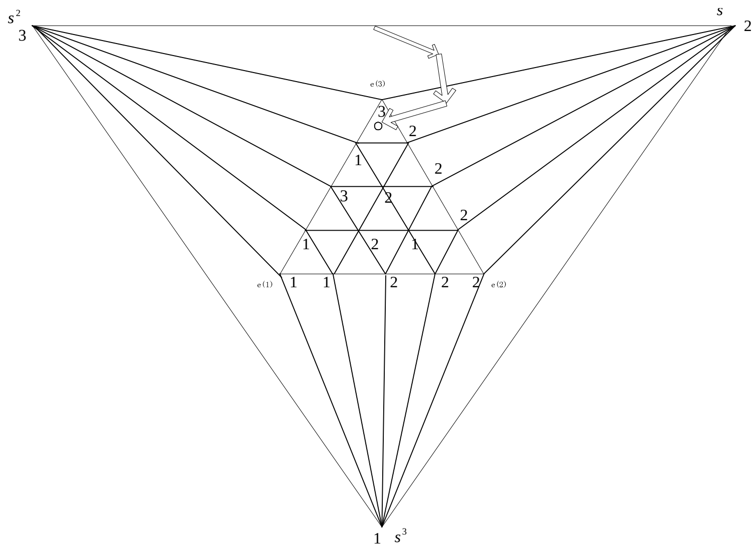

Figure 3.2: The Scarf's triangulation.

We define a proper labeling rule: $l : l(p) = s_{ij}$, where (i, j) is the lexicographic least index in $\underset{i \in N, 1 \le j \le m_i}{\arg\max} f_{ij} - p_{ij}$ [66, 93]. We apply these transformations to the above algorithm to compute the Nash equilibria.

3.3.3 Other Algorithms

In addition to Scarf's algorithms, there are many other algorithms for solving fixed point problems. In 1968, Kuhn introduced the first simplicial subdivision algorithm. Kuhn's artificial start algorithm uses the $K_2(m)$-triangulation without imbedding the $K_2(m)$-triangulation into the original triangulation, but introduced an extra layer below the unit triangulation [113]. It also uses a different labeling rule

$$l(x) = \min\{j \in I_n | x_j = \max_h x_h\}$$

where x_j is defined as in the $K_2(m)$-triangulation. In the following year, Kuhn introduced his second fixed point algorithm, which used the unit $K_2(m)$-triangulation, and starts from one of the vertices of the unit simplex. This algorithm allows the dimension of adjacent simplices to change in the process of tracing. Some other simplicial subdivision algorithms were developed by Merrill (1972), Kuhn and MacKinnon (1975), van Der Lann and Talman (1979, 1980, 1982), van der Laan, Talman and van der Heyden (1987), and Doup and Talman (1987) [66, 113]. These algorithms are reported to be more computationally efficient.

Another branch of algorithms for finding fixed points is called simplicial homotopy. There algorithms were first introduced by Eaves (1972), Eaves and Saigal (1972), and Merrill(1972). This algorithm is still a triangulation algorithm, but can start at anywhere and automatically refines the grid size of simplices in the process [113]. [2]

3.3.4 Complexity Results

The simplicial subdivision algorithms approximate Nash equilibria. The computational efficiency depends on the triangulation expanding method used by the specific algorithm, the grid size of a single simplex, and the precision required. Although there is still room for further improvement, so far, all algorithms are NP-hard. It has been shown by Hirsch that the worst case running time to compute a Brouwer fixed point is exponential in the size of the triangulation and the accuracy we require [33]. Saigal shows that a proper alternation of the mesh size of the simplices can improve the convergence quicker [91]; however, there is no algorithm can easily solve an n-person Nash equilibrium mapping from fixed point for a large size game in general.

3.4 Extensive Form Games

The algorithms presented above are for normal form games. As we know, games in extensive form may contain more information than those in normal form. Generally, when we reduce an extensive form game to normal form games, we lose information about the sequence of moves. Moreover, there are several schemes that we may use to reduce the extensive form games to normal form games. The most popular ones are multiplication of behavior strategy, agent normal form, and reduced normal form. As discussed in Chapter 2, different reductions can lead to different solutions. Although we may argue that the game in nature is the same, the computation procedure and the complexity of the computing equilibria will depend on the game form we use. Regardless of which reduction scheme they choose, in the worst case computation is still exponential in the size of the behavior strategy in extensive form.

[2] Besides the researchers mentioned above, there are still many others contributed to this area. Please refer to [113] and the other references for more details.

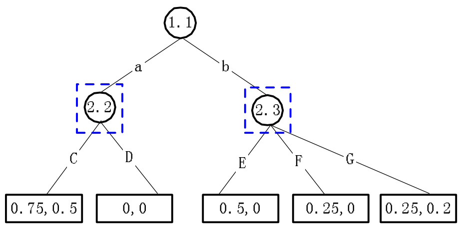

Figure 3.3: Illustration of sequence form.

Nevertheless, do we have to reduce extensive forms to normal form in order to compute the games? In fact, some papers point out that another game form, *sequence form*, will enable us to solve the extensive form game more efficiently [44, 45, 95, 104, 105].

Definition 3.4.1. *A representation of an extensive form game is called* sequence form *if it is defined by*

1. *A empty sequence $\emptyset$ for each player. Nature will be regarded as Player 0.*

2. *A* sequence *is a possible choice for player i from the root to a specified node. The set of sequences of Player i is denoted by S_i.*

3. *The payoff of Player i is defined as the payoff combination of sequences of Player i.*

4. *the probability of sequence $s_i \in S_i$ is defined by a* realization play

$$r_i(s_i) = \prod_{c \in s_i} s_i(c)$$

where s_i is the behavior strategy of player i.

Let us look at an example as shown in Figure 3.3. The sequences of player 1 is a, b, and $\emptyset$, the empty sequence. The sequences of player 2 are C, D, E, F, G, and $\emptyset$, the empty sequence. However, in normal form, the possible strategies for player 2 are CE, CF, CG, DE, DF, and DG. The advantage of sequence form lies in that the size of sequences is linear in the size of extensive form, while any other reduction game form might still be exponential in the size of the game tree. However, like every coin has two sides, the definition of a realization plan is in sequence form less intuitive than behavior strategy probability [104].

3.5 Finding All Equilibria

The multiplicity of Nash equilibria not only hinders the forecasting capability, but also complicates the computability. The algorithms for computing one Nash equilibrium cannot guarantee that the result has salient features. For example, the identified equilibrium might be Pareto dominated. Thus, finding a single equilibrium may not always satisfy us.

However, it is very expensive and sometimes impossible to compute all equilibria for a game. McLennan [70, 71] derived several results on computing the expected number of equilibria. In general, however, it is very difficult to predict the exact number of equilibria for a game, and in some cases, the number of equilibria of a game is infinite.

One new approach that sounds promising is semi-algebraic set algorithms. The Nash equilibrium can be expressed as a conjunction of polynomial equations and weak inequalities [66]. Finding all equilibria is equivalent to finding all roots of the equations and inequalities.

3.6 Summary

Computing equilibria is still a fertile research area although there are many algorithms already developed. Table 3.1 is a summary of algorithms for finding a single Nash equilibrium or its refinements. Table 3.2 summarizes the algorithms for computing all equilibria.

Unfortunately, current algorithms can solve only small sized games. The following is a summary of current research on computing equilibria.

1. Algorithms for computing 2-person games. Although this category of games is the easiest one, many important results are based on 2-person games, including the minimax theorem and the Lemke-Howson algorithm. Moreover, algorithms for 2-person games are easier to understand intuitively. There are still many researchers active in this area.

2. Algorithms for computing n-person games. These approaches normally seek a mapping from a Nash equilibrium to a mathematical problem, such as fixed point problem, stationary point

Equilibrium	2-person normal form games	2-person extensive form games	n-person normal form games	n-person extensive form games
Nash	Lemke-Howson Algorithm (LH 64)	LH Variant (Wilson 72), LH Variant (Koller, etc. 94, Knowledge Representation)	Simplicial Subdivision (Scarf 67,73);LH Variant (Rosenm$\ddot{u}$ller 71,Wilson 71); Nonlinear Complementarity Problem (Mathiesen 87); Minimum Method (McKelvey 92); Global Newton Method (Govindan&Wilson 98)	
Bayesian				
Subgame Perfect				
Perfect Bayesian	Lex-Order LH Variant (Eaves 71)	Homotopy Method (von Stengel,etc,2002)		
Sequential				Homotopy Based Algorithm (McKelvey & Palfrey 94);B-labeling(Azar. etc. 2000)
Persistent				
Proper			Homotopy Method(Yamamoto 93); Simplicial Subdivision (Talman&Yang 94)	
Stable	Lex-Order LH (Wilson 92)		Exhaustive Triangulation (Mertens 88, an idea only)	

Table 3.1: Algorithms to find a single equilibrium.

Equilibrium	2-person normal	2-person extensive	n-person normal	n-person extensive
Nash	LcpSolve (Gambit) [67]	Koller and Megiddo 94	PolEnumSolve (Gambit) [68]; QreSolve, McKelvey and Palfrey 95 [68][4]	Lyapunov function method, McKelvey 91 [65][4]
Correlated		von Stengel 2001		
Bayesian				
Subgame Perfect				
Perfect Bayesian				
Sequential				QreSolve, McKelvey and Palfrey 98 [69]
Proper				
Persistent				
Stable	Lex-Order LH (Wilson 92)			

Table 3.2: Algorithms to compute all equilibria.

problem, etc. The mathematics literature focuses on solving the math problem and thus provides a relevant algorithm for computing equilibria.

3. Algorithms for computing a special class of games, such as auction games. For some subclasses of games, the computational cost can be reduced dramatically since we can take advantage of special structure in the problem. An example is the task to compute the equilibrium of FPSB auctions in the following chapter.

[4]This algorithm can find multiple equilibria for both normal form and extensive form games; however, it is not guaranteed to find all equilibria [67].

Chapter 4

Equilibrium Strategies in First-Price Sealed-Bid Auctions with Discrete Bids

4.1 Introduction

In a *first-price sealed-bid (FPSB) auction*, each agent submits a single bid without observing others' bids, and the agent with the highest bid pays the value of its bid. In terms of information and bidding space, FPSB auctions can be categorized into four subclasses: FPSB auctions with incomplete information and continuous bids; FPSB auctions with incomplete information and discrete bids; FPSB auctions with complete information and continuous bids; FPSB auctions with complete information and discrete bids.

In the past decades, a voluminous theoretical literature has been developed on FPSB auctions. A class of typical models assume that agents are symmetric and risk neutral [62, 73, 79, 87, 102]. More recently, researchers have studied models with asymmetric information [16, 18, 54, 59, 60], affiliated values [1][11, 61], single-crossing[2][2], or other restrictions [51, 55, 57, 58]. Most of the literature assumes incomplete information and continuous bids. With agents having complete information and continuous bids, equilibrium does not exist due to the discontinuity of the payoff

[1]A condition is *affiliated* if higher values of some variables make higher values of the others more likely [73, 106].
[2]The *single crossing condition* holds when agents with higher types will choose higher strategies [2].

function [53]. To enable the existence of equilibrium in this situation, Lebrun [53] suggests an "augmented" first price auction to break the tie so that the agent having higher valuation wins the item by submitting a "message" marking its higher value. Maskin and Riley [60] propose using a second round Vickrey auction to break the tie if any. In a model of incomplete information, Athey [2] concludes the existence of pure strategy equilibrium in FPSB with finite strategies, when the single crossing condition holds. However, the properties of existence and multiplicity of equilibria in FPSB auctions with discrete bids, especially the relationship of multiplicity of equilibria to the size of bid increment, are not well studied. Discrete valued bidding is relatively common in online auctions with fixed ending times such that they degenerate to FPSB auctions.[3]

When the number of discrete bids is finite or bound, the existence of equilibrium in FPSB auctions with complete information can be easily proven using the Nash theorem [81, 83]. In this chapter, we discuss the existence of equilibrium from a different perspective and focus on the multiplicity of equilibria in FPSB auctions with complete information. We also discuss the equilibria in FPSB auctions with incomplete information. The remainder of this chapter is organized as follows. In Section 4.2, we present a model of the FPSB auction. In Section 4.3, we discuss the multiplicity of equilibria in two-agent FPSB auctions and multi-person FPSB auctions, including sequential FPSB auctions, with complete information. In Section 4.4, we provide equilibrium solutions for both two-person and multi-person FPSB auctions with incomplete information. We offer some concluding remarks in Section 4.5.

4.2 The Model

Assume that there is an item for sale in an FPSB auction. There are n agents competing for this item. The set of agents is denoted A. In this chapter, we discuss two different scenarios: one is $n = 2$; another one is $n > 2$.

We assume that these agents are risk-neutral and the payoff of each agent is equal to its monetary surplus. We adopt a random tie-breaking rule in which each agent has $1/\#(tie)$ probability to win the item, an assumption that is common in research. Let u_i denote the utility of agent i.

In this chapter, we explore both FPSB auctions with complete information and FPSB auctions with incomplete information. Complete information means that the payoff functions are

[3]The minimum bidding unit is one cent in most online auctions.

visible to all agents. Let v_i be the true valuation of agent i. To simplify the discussion, we assume $v_1 \leq ... \leq v_i \leq v_{i+1} \leq ... \leq v_n$. In incomplete information cases, we continue to let v_i be the true valuation of agent i. Each agent knows the value of the object to itself, and the valuation of different bidders are independent observations of a nonnegative random variable, Y, from a commonly known continuous distribution, F, and its associated probability density function, f. Let Y_j denote the valuation variable of agent j. Without loss of generality, let Y_j be the $(n - j + 1)$-st order statistic of $\{Y_1, ..., Y_{n-1}\}$. Thus, we have $Y_1 \leq Y_2 \leq ... \leq Y_{n-1}$.

We assume that agents have the same discrete strategy (bidding) space, denoted by B. Let δ be a minimum possible bid increment in the auction, which is the smallest value between any two bids and is positive and fixed. The difference between bids can be as small in the real world as, for example, one cent. Let b_k be the k^{th} element in B and b_0 denote the lowest bid value in B. Without loss of generality, let B be ordered such that $b_k + \delta = b_{k+1}$.

In complete information cases, let s_i, $s_i \in B$, denote the strategy that agent i uses in this auction. In incomplete information cases, let $\beta(v_i)$ denote the strategy function of agent i given agent i has true valuation v_i. We assume that $s_i \leq v_i$ and $\beta(v_i) \leq v_i$, which means that agents will not bid higher than their own valuations.

We define $\lfloor v_i \rfloor \equiv [(v_i - b_0) \mod \delta] * \delta + b_0$. It is easy to see that $\lfloor v_i \rfloor$ is the highest possible bid value for agent i because agent i will get negative surplus if it bids $\lfloor v_i \rfloor + \delta$ and wins. Whether $b = \lfloor v_i \rfloor$ is an equilibrium strategy for agent i depends on the situation. In a two-agent FPSB auction, for example, if v_2 is larger than v_1, say $v_2 = v_1 + 10\delta$, $\{\lfloor v_1 \rfloor, \lfloor v_1 \rfloor + \delta\}$ is an equilibrium. In contrast, if $v_2 = v_1$, $\{\lfloor v_1 \rfloor - \delta, \lfloor v_1 \rfloor - \delta\}$ is an equilibrium. Thus, to bid at $\lfloor v_1 \rfloor$ is not the only equilibrium strategy for agent 1. The agents need to judge whether or not to tie when a tied strategy could be a Nash equilibrium.

We call a strategy profile an *identical bid profile* if all agents bid at the same price. If an identical bid profile is an equilibrium, we call it an *identical bid equilibrium*. As it turns out, it is helpful to discuss the equilibrium in the FPSB auction in terms of identical bid profiles.

4.3 FPSB Auctions with Complete Information

4.3.1 A Two-Person FPSB Auction

We start from a two-person FPSB auction. The utility function for agent 1 is

$$u_1 = \begin{cases} v_1 - s_1, \text{ if } s_1 > s_2 \\ \frac{1}{2}(v_1 - s_1), \text{ if } s_1 = s_2 \\ 0, \text{ if } s_1 < s_2. \end{cases} \tag{4.1}$$

The utility function for agent 2 is

$$u_2 = \begin{cases} v_2 - s_2, \text{ if } s_2 > s_1 \\ \frac{1}{2}(v_2 - s_2), \text{ if } s_2 = s_1 \\ 0, \text{ if } s_2 < s_1. \end{cases} \tag{4.2}$$

Suppose that both agents use an identical bid. If no agent can benefit from deviating unilaterally, this identical bid profile is a Nash equilibrium. If some agent would be better off by deviating unilaterally, the only possibly positive deviation is to bid higher. Thus, an identical bid profile $\{b, b\}$ is a Nash equilibrium if and only if

$$\begin{aligned} & s_1 = s_2 = b, b \in B, \\ & \tfrac{1}{2}(v_i - b) \geq v_i - b - \delta, i = 1, 2, \\ & b \leq \min_i \{v_i\}, \end{aligned} \tag{4.3}$$

where the second condition implies that no agent can be better off by deviating unilaterally. The third condition requires no agent bids higher than its true valuation, which is rational for agents.

Lemma 4.3.1. *For all $b \in B$, $\{b, b\}$ is a pure strategy equilibrium in a 2-person FPSB auction with complete information and a discrete bidding space if and only if*

$$\max_i \{v_i\} - 2\delta \leq b \leq \min_i \{v_i\}, i = 1, 2.$$

Proof: From the second condition in equation (4.3), we have $b \geq (v_i - 2\delta)$, for $i = 1, 2$. Because it is true for all agents, $b \geq \max_i\{v_i\} - 2\delta$. When we combine this with the third constraint in equation (4.3), we have $\max_i\{v_i\} - 2\delta \leq b \leq \min_i\{v_i\}$. $\diamond$

Lemma 4.3.1 suggests that we need to check the conditions in equation (4.3) for every b. In fact, we do not have to check every element in the strategy space. We can constrain the range of possible equilibria by finding the lowest bid value b_c such that $\{b_c, b_c\}$ is a Nash equilibrium. From Lemma 4.3.1, we have $\max_i\{v_i\} - 2\delta \leq b_c \leq \min_i\{v_i\}$. Because b_c is the lowest value in the above range, it requires $b_c - (\max_i\{v_i\} - 2\delta) < \delta$. Combining these two equations, we have

$$0 < b_c - (\max_i\{v_i\} - 2\delta) < \delta.$$

We call $\{b_c, b_c\}$ a *critical identical bid equilibrium*, if it exists. It is easy to prove that $\{b_c, b_c\}$ is a weakly Pareto dominant equilibrium when there are more identical bid equilibria.

Lemma 4.3.2. *In a two-person FPSB auction with complete information and a discrete bidding space B, if there exists a critical identical bid equilibrium $\{b_c, b_c\}$, then $\{s, s\}$ is also a pure strategy equilibrium, if and only if for all $s \in B$*

$$b_c < s \leq \min_i\{v_i\}, i = 1, 2.$$

Proof: From the definition of the critical equilibrium, it is obvious that no agent can unilaterally bid lower than b_c and be better off. So, the only possibility is that some agent would have an incentive to bid higher. To be an identical equilibrium, $s \leq \min_i\{v_i\}$. Combining $\max_i\{v_i\} - 2\delta \leq b_c$ and $b_c < s$, we have $\max_i\{v_i\} - 2\delta \leq s$. Thus,

$$\max_i\{v_i\} - 2\delta \leq s \leq \min_i\{v_i\}, i = 1, 2.$$

By Lemma 4.3.1, we know that $\{s, s\}$ is a pure strategy equilibrium. $\diamond$

Both Lemma 4.3.1 and Lemma 4.3.2 are supported by Example 4.3.1.

Example 4.3.1. *Let $v_1 = 6.5$ and $v_2 = 7$. Suppose that $\delta = 1$ and $b_0 = 0$. There are two equilibria in pure strategies. One is $\{6, 6\}$. Another one is $\{5, 5\}$.*

Combined, Lemmas 4.3.1 and 4.3.2 give us a partial picture of the properties of identical bid equilibria in a two-person FPSB auction. We now establish the conditions under which non-identical bid equilibria exist. We let $b_k = \delta \times ((\min_i\{v_i\} - b_0) \mod \delta) + b_0$, and $b_{k+1} = b_k + \delta$; $b_k, b_{k+1} \in B$. The condition of existence of a non-identical bid equilibrium can be obtained by the following.

Lemma 4.3.3. *In a two-person FPSB auction with complete information and a discrete bidding space, a non-identical bid profile $\{b_k, b_{k+1}\}$ is a Nash equilibrium if and only if*

$$\delta < \frac{1}{2}(\max_i\{v_i\} - b_k).$$

When $b_k < \min_i\{v_i\} < b_{k+1}$, $\{b_k, b_{k+1}\}$ is also unique.

Proof: The definition of $b_k = \delta \times ((\min_i\{v_i\} - b_0) \mod \delta) + b_0$ tells us that b_k is the highest possible bid with non-negative surplus for the lower type agent. From equation (4.3), an identical bid profile $\{b_k, b_k\}$ could be an equilibrium if and only if $\delta \geq \frac{1}{2}(v_i - b_k), i = 1, 2$. $\delta < \frac{1}{2}(\max_i\{v_i\} - b_k)$, however, implies that the higher type agent will be better off if it deviates from an identical bid profile $\{b_k, b_k\}$. Because

$$\begin{aligned} v_2 - b_{k+1} &= v_2 - (b_k + \delta) \\ &= v_2 - b_k - \delta \\ &> v_2 - b_k - \frac{1}{2}(v_2 - b_k) \\ &= \frac{1}{2}(v_2 - b_k). \end{aligned}$$

When the higher type agent deviates to b_{k+1}, it is an equilibrium strategy for the lower type agent to bid b_k, because it cannot be better off by bidding less or more than b_k since b_k is its highest affordable bid. Meanwhile, it would be irrational for the higher type agent to bid b_{k+2} or even higher, if it can win with b_{k+1}. Thus, $\{b_k, b_{k+1}\}$ must be an equilibrium because no agent wants to deviate.

From Lemma 4.3.2, we know that if $\{b_k, b_k\}$ is not an equilibrium, no other identical bid profile can be an equilibrium.

On the other hand, a strategy profile cannot be an equilibrium if the lower type agent bid lower than b_k. Because the higher type agent will be better off by deviating from b_{k+1} to b_k;

however, the lower type agent will switch back to b_k with an attempt to tie with the higher type agent to gain a positive surplus when $b_k < \min_i\{v_i\} < b_{k+1}$. Thus, if $\delta < \frac{1}{2}(\max_i\{v_i\} - b_k)$ and $b_k < \min_i\{v_i\} < b_{k+1}$, $\{b_k, b_{k+1}\}$ will be unique. $\diamond$

However, if the condition $\delta < \frac{1}{2}(\max_i\{v_i\} - b_k)$ does not hold, $\{b_k, b_{k+1}\}$ is not guaranteed to be an equilibrium. The following lemma provides a relationship between the identical bid equilibrium and the non-identical bid equilibrium.

Lemma 4.3.4. *In a two-person FPSB auction with complete information and a discrete bidding space, if $\{b_k, b_{k+1}\}$ is not an equilibrium, then $\{b_k, b_k\}$ must be an equilibrium.*

Proof: $b_k = \delta \times ((\min_i\{v_i\} - b_0) \mod \delta) + b_0$ is the highest bid the lower valuation agent can play; otherwise it will get negative payoff. The agent with the higher valuation does not want to deviate to $b_{k+n}, n \geq 2$, if it can win at b_{k+1}. Thus, the reason that the higher type agent does not want to play b_{k+1} is that it can be better off by bidding lower. So, the higher type agent will deviate from b_{k+1} to b_k. At $\{b_k, b_k\}$, the higher type agent will not bid lower because it will obtain nothing if it bids lower than b_k. For the lower type agent, it will not bid lower or higher than b_k because it cannot be better off by doing so. As a result, neither agent wants to deviate unilaterally from $\{b_k, b_k\}$. So, $\{b_k, b_k\}$ must be an equilibrium if $\{b_k, b_{k+1}\}$ is not an equilibrium. $\diamond$

Theorem 4.3.5. *In a two-person FPSB auction with complete information and a discrete bidding space, there exists at least one equilibrium.*

Proof: If $\delta < \frac{1}{2}(\max_i\{v_i\} - b_k)$, from Lemma 4.3.3, we know that there is one unique equilibrium $\{b_k, b_{k+1}\}$. Otherwise, from Lemma 4.3.1, 4.3.2, and 4.3.4, we know that there is at least one identical bid equilibrium. $\diamond$

So far, we have not addressed the issue of how many equilibria exist in a two-person FPSB auction with complete information and a discrete bidding space. Since $b_k = \delta \times ((\min_i\{v_i\} - b_0) \mod \delta) + b_0$, we have $b_k \leq \min_i\{v_i\} < b_{k+1}$. This condition holds in the following discussion unless otherwise mentioned. From the above results, we may categorize all possible equilibrium results into eleven situations, as shown in Figure 4.1. Specifically, we have:

1. **Situation A**: When $b_k < \max_i\{v_i\} - 2\delta$ and $b_k < \min_i\{v_i\} < b_{k+1}$, $\{b_k, b_{k+1}\}$ is the unique equilibrium.

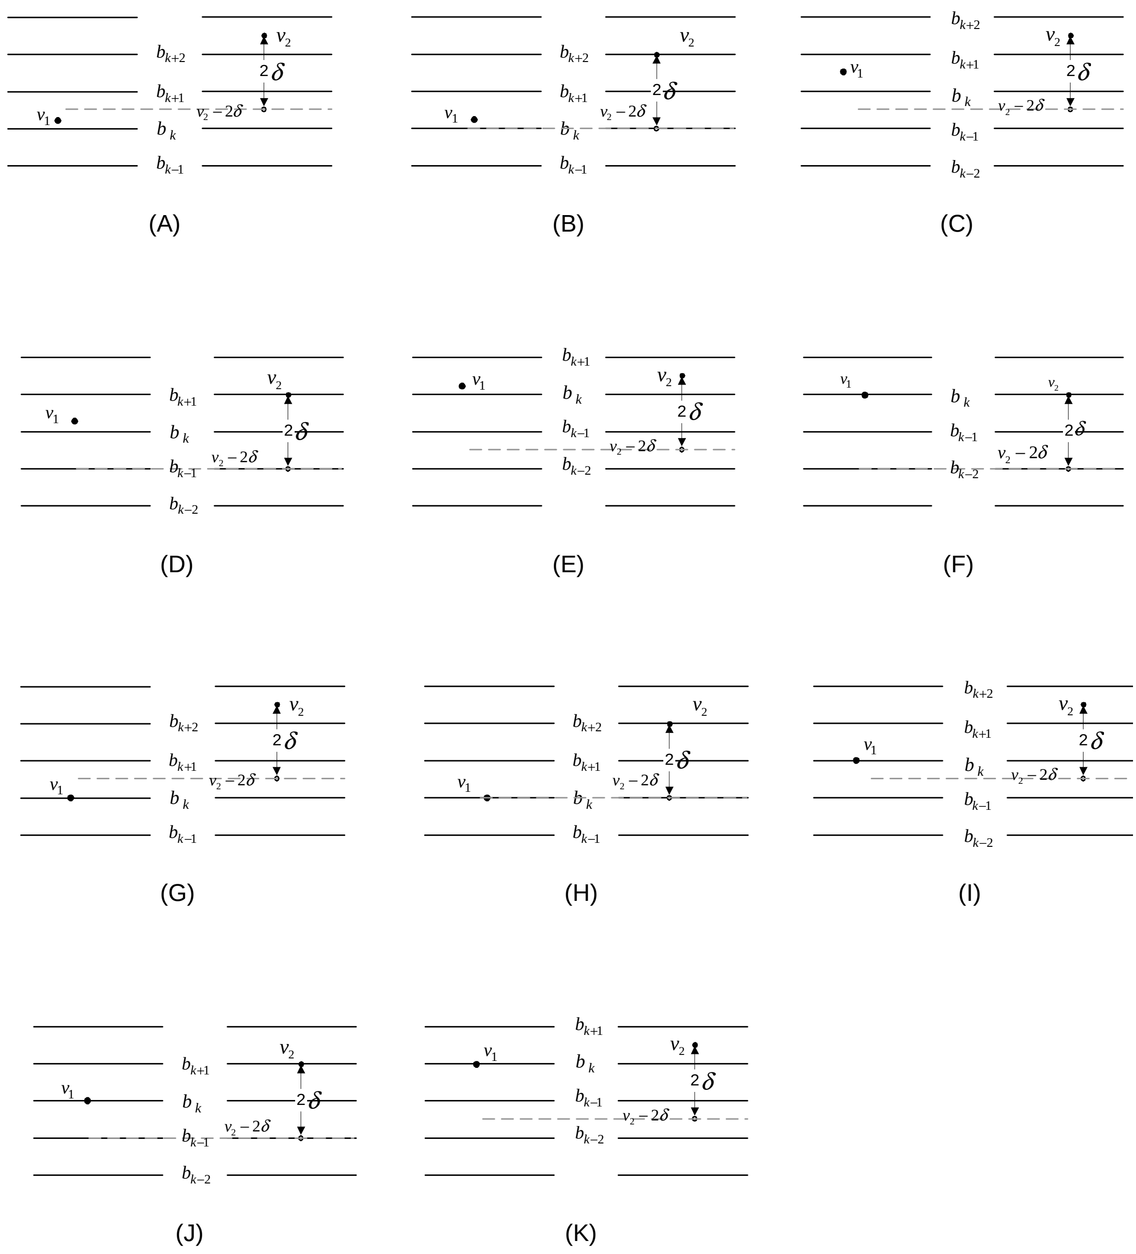

Figure 4.1: Eleven possible equilibrium situations, given that $v_1 = \min_i\{v_i\}$ and $v_2 = \max_i\{v_i\}$.

2. **Situation B**: When $b_k = \max_i\{v_i\} - 2\delta$ and $b_k < \min_i\{v_i\} < b_{k+1}$, $\{b_k, b_k\}$ and $\{b_k, b_{k+1}\}$ are the only two equilibria.

3. **Situation C**: When $\max_i\{v_i\} - 2\delta < b_k < \max_i\{v_i\} - \delta$ and $b_k < \min_i\{v_i\} < b_{k+1}$, $\{b_k, b_k\}$ is the unique equilibrium.

4. **Situation D**: When $b_{k-1} = \max_i\{v_i\} - 2\delta$ or $b_k = \max_i\{v_i\} - \delta$, and $b_k < \min_i\{v_i\} < b_{k+1}$, $\{b_{k-1}, b_{k-1}\}$ and $\{b_k, b_k\}$ are the only two equilibria.

5. **Situation E**: When $\max_i\{v_i\} - 2\delta < b_{k-1} < \max_i\{v_i\} - \delta$ or $\max_i\{v_i\} - \delta < b_k < \max_i\{v_i\}$, and $b_k < \min_i\{v_i\} < b_{k+1}$, $\{b_{k-1}, b_{k-1}\}$ and $\{b_k, b_k\}$ are the only two equilibria.

6. **Situation F**: When $b_{k-2} = \max_i\{v_i\} - 2\delta = \min_i\{v_i\} - 2\delta$ or $b_k = \max_i\{v_i\} = \min_i\{v_i\}$, $\{b_{k-2}, b_{k-2}\}$, $\{b_{k-1}, b_{k-1}\}$, and $\{b_k, b_k\}$ are the only three equilibria.

7. **Situation G**: When $b_k < \max_i\{v_i\} - 2\delta$ and $b_k = \min_i\{v_i\}$, $\{b_{k-1}, b_k\}$ and $\{b_k, b_{k+1}\}$ are the only two equilibria.

8. **Situation H**: When $b_k = \max_i\{v_i\} - 2\delta$ and $b_k = \min_i\{v_i\}$, $\{b_{k-1}, b_k\}$ and $\{b_k, b_k\}$ and $\{b_k, b_{k+1}\}$ are the only three equilibria.

9. **Situation I**: When $\max_i\{v_i\} - 2\delta < b_k < \max_i\{v_i\} - \delta$ and $b_k = \min_i\{v_i\}$, $\{b_{k-1}, b_k\}$ and $\{b_k, b_k\}$ are the only two equilibria.

10. **Situation J**: When $b_{k-1} = \max_i\{v_i\} - 2\delta$ or $b_k = \max_i\{v_i\} - \delta$, and $b_k = \min_i\{v_i\}$, $\{b_{k-1}, b_k\}$, $\{b_{k-1}, b_{k-1}\}$ and $\{b_k, b_k\}$ are the only three equilibria.

11. **Situation K**: When $\max_i\{v_i\} - 2\delta < b_{k-1} < \max_i\{v_i\} - \delta$ or $\max_i\{v_i\} - \delta < b_k < \max_i\{v_i\}$, and $b_k = \min_i\{v_i\}$, $\{b_{k-1}, b_{k-1}\}$ and $\{b_k, b_k\}$ are the only two equilibria.

From Situation A to Situation K, notice that the distance between b_k and v_2 changes from larger than 2δ to 0 when $b_k \leq \min_i\{v_i\} < b_{k+1}$, which implies that these eleven situations are exhaustive and include all possible equilibrium situations in a two-person FPSB auction with complete information and a discrete bidding space.

Theorem 4.3.6. *In a two-person FPSB auction with complete information and a discrete bidding space, there are at least one and at most three equilibria. The concrete situations are given from Situation A to Situation K.*

4.3.2 An n-person FPSB Auction

The results for two-person auctions can be extended to n-person auctions. Again, we start from the identical bid equilibrium. In FPSB auction, a strategy profile is a *k-identical bid profile* if the top highest k, $2 \leq k \leq n$, type agents bid at the same price while others bid at $\lfloor v_i \rfloor$. If a k-identical bid profile is an equilibrium candidate, the top k agents should be willing to tie at a price higher than v_{n-k}. Normally, an agent with a higher valuation will have more power in the negotiation. If there is any agent that can be better off by deviating unilaterally, the agent with the highest valuation should be the first one to do that. However, if the highest type agent cannot be better off by bidding higher, no other agent can be better off by bidding higher. Thus, a k-identical bid profile $\{\lfloor v_1 \rfloor, ..., \lfloor v_{n-k+1} \rfloor, b, ..., b\}$ is a Nash equilibrium if and only if for all $j = n-k+1, ..., n$, for all $l = 1, ..., n-k$ and for all $2 \leq k \leq n$

$$\begin{aligned} s_j &= b, b \in B, \\ s_l &= \lfloor v_l \rfloor < b, \\ \frac{1}{k}(v_j - b) &\geq v_j - b - \delta, \\ b_j &\leq \min_j \{v_j\}. \end{aligned} \tag{4.4}$$

The first two conditions state the strategies agents use in the profile. $\lfloor v_l \rfloor < b$ in the second condition requires b larger than the true valuations of all lower type agents. The third condition requires that no agent can be better off by bidding higher. The fourth condition requires an agent bid no higher than its true valuation. Because we can express and solve this problem using linear programming, the complexity of computing an equilibrium of this auction is polynomial. The definition of k-identical bid profile in equation (4.4) leads to the following conclusion.

Lemma 4.3.7. *In an n-person FPSB auction with complete information and a discrete bidding space, a k-identical bid profile $\{\lfloor v_1 \rfloor, ..., \lfloor v_{n-k+1} \rfloor, b, ..., b\}$ is an equilibrium if and only if there exists a b such that for all $j = n-k+1, ..., n$*

$$\max_j \{v_j\} - \frac{k}{k-1}\delta \leq b \leq \min_j \{v_j\}.$$

Proof: The right hand side of the equation is directly given by the fourth condition of equation (4.4). From the third condition of equation (4.4), we have $v_j - \frac{k}{k-1}\delta \le b$. Given that the highest type agent does not want to deviate unilaterally, we have $\max_j\{v_j\} - \frac{k}{k-1}\delta \le b$. $\diamond$

Lemma 4.3.8. *In an n-person FPSB auction with complete information and a discrete bidding space, if there is not a k-identical bid equilibrium, for all $k \ge 2$, $\{\lfloor v_1 \rfloor, ..., \lfloor v_{n-2} \rfloor, \lfloor v_{n-1} \rfloor, \lfloor v_{n-1} \rfloor + \delta\}$ must be a non-identical bid equilibrium.*

Proof: If there is not a k-identical bid equilibrium, it implies that there does not exist a $b \in B$ that satisfies all the constraints in equation (4.4). In other words, there is no agent that wants to tie with other agents. In particular, the agent with the highest valuation does not want to tie with any lower type agents. Since there is no k-identical bid equilibrium, for all $k \ge 2$, the strategy profile $\{\lfloor v_1 \rfloor, ..., \lfloor v_{n-1} \rfloor, \lfloor v_{n-1} \rfloor + \delta\}$ is a potential equilibrium. In this strategy profile, no agent will deviate to a lower bid. At the same time, we know that all agents, except the agent with the highest valuation, have already chosen their highest acceptable strategies. The only possibility is that the agent with the highest valuation deviates to a higher bid and we obtain a new strategy profile; however, the agent with the highest valuation would not bid any higher because it can win the auction with $\lfloor v_{n-1} \rfloor + \delta$. As a result, no agent can be better off by deviating unilaterally. Thus, $\{\lfloor v_1 \rfloor, ..., \lfloor v_{n-1} \rfloor, \lfloor v_{n-1} \rfloor + \delta\}$ is an equilibrium. $\diamond$

Like analysis of the two-agent case, it is worth noting that the equilibria in Lemmas 4.3.7 and 4.3.8 are not unique. In an FPSB auction with complete information, the results are determined by whether the top type agent would like to tie with the others. For a k-identical bid equilibrium in Lemma 4.3.7, it is optional for the bottom $n-k$ agents to bid any bids $\le \lfloor v_i \rfloor, i \le n-k$. Similarly, with the bottom $n-2$ agents bidding any bids $\le \lfloor v_i \rfloor, i \le n-2$, and top two agents bidding at $\lfloor v_{n-1} \rfloor + \delta$ and $\lfloor v_{n-1} \rfloor$ respectively, the new strategy profile still constructs an equilibrium in Lemma 4.3.8.

Through similar reasoning as in a 2-person FPSB auction, we conclude:

Theorem 4.3.9. *In an n-person FPSB auction with complete information and a discrete bidding space, there exists at least one equilibrium.*

4.3.3 Equilibrium in Sequential FPSB Auctions

In this section, we extend the discussion from one individual FPSB auction to sequential FPSB auctions, in which there are K items for sale in K individual FPSB auctions.

It is easy to understand that an agent with a higher valuation will have more power in negotiation and competition. For example, if agent i achieves surplus x from the sequential auctions, agent $i+1$ can also get at least x from the same game by simply taking the same strategy as agent i because $v_i \leq v_{i+1}$. Hence, we may expect that in sequential FPSB auctions with complete information, an agent with higher valuation will yield no less surplus than a lower type agent.

Theorem 4.3.10. *In sequential FPSB auctions with complete information and a discrete bidding space, if* $v_{n-K+1} > \lfloor v_{n-K} \rfloor + 2\delta$*, then* $\{\lfloor v_1 \rfloor, ..., \lfloor v_{n-K} \rfloor, \lfloor v_{n-K} \rfloor + \delta, ..., \lfloor v_{n-K} \rfloor + \delta\}$ *is an equilibrium for all agents in every single FPSB auction.*

Proof: Let us consider the last auction. Suppose in the first $K-1$ auctions, $K-1$ of the highest K valuation agents win and leave. From Lemma 4.3.3 and 4.3.7, agent $n-K+1$ can be better off by deviating from an identical bid profile in which she has the same bid as agent $n-K$, if condition $v_{n-K+1} > \lfloor v_{n-K} \rfloor + 2\delta$ is true. Thus $\lfloor v_{n-K} \rfloor + \delta$ is an equilibrium strategy for agent $n-K+1$ in the last auction.

Now let us discuss the first $K-1$ auctions. In a complete information sequential auctions, an agent with higher valuation will yield no less surplus than an agent with lower valuation. Thus, in the $(K-1)st$ auction, no agent wants to bid higher than $\lfloor v_{n-K} \rfloor + \delta$, because, if it lost in the $K-1$ auction, it could bid $\lfloor v_{n-K} \rfloor + \delta$ to win the last auction. By induction, no agent wants to bid more than $\lfloor v_{n-K} \rfloor + \delta$ in the whole game. $\diamond$

The above equilibrium is not unique in sequential FPSB auctions with complete information and a discrete bidding space, given that $v_{n-K+1} \leq \lfloor v_{n-K} \rfloor + 2\delta$. For example, in the final auction, if agent $n-K+1$ can benefit from an identical bid with $n-K$ agents or more agents, then agent $n-K+1$ might use a k-identical bid profile while the top $K-1$ type agents bid at $\lfloor v_{n-K} \rfloor + \delta$ in the first $K-1$ auctions.

4.3.4 Equilibrium in Multi-unit Sequential Auctions

It is not surprising that some of the above conclusions for the FPSB sequential auctions will hold for sequential multi-object auctions, in which a single auction consists of multiple units, but agents demand a single-unit demand. Given some single auctions are pay-your-bid auctions, the results will be the same as in the FPSB sequential single-object auctions. Because agents pay what they bid, no agent wants to bid more than $\lfloor v_{n-K} \rfloor + \delta$ if the condition $v_{n-K+1} > \lfloor v_{n-K} \rfloor + 2\delta$ holds. If $v_{n-K+1} \leq \lfloor v_{n-K} \rfloor + 2\delta$, the solution will also be exactly the same as in sequential FPSB auctions. Similarly, in sequential *M*th price auctions, no agent wants to bid more than $\lfloor v_{n-K} \rfloor + \delta$, given that $v_{n-K+1} > \lfloor v_{n-K} \rfloor + 2\delta$ is true.

4.4 FPSB Auctions with Incomplete Information

We now turn our attention to the case where the bidders do not know the valuations of the other bidders.

4.4.1 A Two-Person FPSB Auction

Again, we start from a two-person FPSB auction. In this auction, each agent knows its own valuation; however, she does not know the other agents' valuations. We assume that agents' valuations are independent observations of a commonly know continuous cumulative distribution function, F. In this symmetric FPSB auction, we assume that there exists a symmetric equilibrium. To simplify the discussion, we first assume that each agent has only two possible strategies $\{b_1, b_2\}$. Suppose that there is a critical value, z, which is a value inside the domain of F. Specifically, we assume that an agent uses strategy b_1, when its valuation is lower than z; otherwise, it uses b_2. Thus, we have the following strategy function for agent i.

$$\beta(v_i) = \begin{cases} b_1, & v_i \leq z \\ \\ b_2, & otherwise. \end{cases} \tag{4.5}$$

Theorem 4.4.1. *In a two-person symmetric FPSB auction with discrete bids and incomplete information, there exists a symmetric equilibrium as featured by equation (4.5), where the z value is* *given by*

$$z = b_2 + F(z)(b_2 - b_1). \tag{4.6}$$

Proof: We let $u_i(v_i, z, b_1)$ be the utility of agent i when it bids b_1 and let $u_i(v_i, z|b_2)$ be the utility of agent i when it bids b_2. Without loss of generality, we consider agent 1 only. Since there is only one other agent, Y_1 is the highest valuation of the other agents. When agent 1 bids b_1, and the other agent also bids b_1, resulting in a tie at b_1, the probability that it wins the item is $\Pr(Y_1 \leq z)$. Thus,

$$\begin{aligned} u_1(v_1, z, b_1) &= \Pr(Y_1 \leq z)\tfrac{1}{2}(v_1 - b_1) \\ &= \tfrac{1}{2}F(z)(v_1 - b_1). \end{aligned}$$

When agent 1 bids b_2, there are two possibilities that agent 1 can win the item. First, if $Y_1 \leq z$, agent 1 has probability $\Pr(Y_1 \leq z)$ to win the item. Secondly, if $Y_1 > z$, agent 1 has probability $\Pr(Y_1 > z)$ tie with agent 2. The utilities can be written as

$$\begin{aligned} u_1(v_1, z, b_2) &= \Pr(Y_1 \leq z)(v_1 - b_2) + \Pr(Y_1 > z)\tfrac{1}{2}(v_1 - b_2) \\ &= F(z)(v_1 - b_2) + \tfrac{1}{2}[1 - F(z)](v_1 - b_2) \\ &= \tfrac{1}{2}(v_1 - b_2) + \tfrac{1}{2}F(z)(v_1 - b_2). \end{aligned}$$

Agent 1 will bid b_1 if $u_1(v_1, z, b_1) \geq u_1(v_1, z, b_2)$; otherwise, it bids b_2. Subtracting $u_1(v_1, z, b_2)$ from $u_1(v_1, z, b_1)$, we obtain

$$\begin{aligned} u_1(v_1, z, b_1) - u_1(v_1, z, b_2) &= \frac{1}{2}F(z)(v_1 - b_1) - [\frac{1}{2}(v_1 - b_2) + \frac{1}{2}F(z)(v_1 - b_2)] \\ &= \frac{1}{2}F(z)(b_2 - b_1) - \frac{1}{2}(v_1 - b_2) \\ &= \frac{1}{2}[F(z)(b_2 - b_1) - (v_1 - b_2)]. \end{aligned}$$

The condition that agent 1 prefers b_1 requires that $\frac{1}{2}[F(z)(b_2 - b_1) - (v_1 - b_2)] \geq 0$, which yields

$$v_1 \leq b_2 + F(z)(b_2 - b_1). \tag{4.7}$$

Combining equation (4.7) and equation (4.5) when $i = 1$, we obtain equation (4.6), $z = b_2 + F(z)(b_2 - b_1)$. $\diamond$

The entire process does not require F be differential in the agent's valuation. Thus, this result is applicable to all possible cumulative distribution functions, although it may be difficult to compute z using equation (4.6) for some distributions. More importantly, there might be a multiplicity of z, which results in a multiplicity of symmetric equilibria. We provide a simple example in Example 4.4.1 when F is uniformly distributed.

Example 4.4.1. *When F is uniformly distributed among $\{c, d\}$, we define*

$$F(v) = \begin{cases} 0, & v < c \\ \frac{v-c}{d-c}, & c \leq v \leq d \\ 1, & d < v. \end{cases}$$

First, let us consider the case when $c \leq z \leq d$. By incorporating F into equation (4.6), we obtain the following:

$$z = b_2 + \frac{z-c}{d-c}(b_2 - b_1),$$

which results in

$$z = \frac{b_2 - \left(\frac{b_2-b_1}{d-c} \times c\right)}{1 - \frac{b_2-b_1}{d-c}}.$$

The value of z is feasible as long as the result is less than d and larger than c. Now we look at the case where $z > d$, we have

$$z = b_2 + (b_2 - b_1),$$

which gives us

$$z = 2b_2 - b_1.$$

This z value is feasible as long as it is larger than d. The third case occurs when $z < c$. This situation gives us $z = b_2$, which requires $b_2 < c$.

4.4.2 A Multi-Person FPSB Auction

The above result can be extended to multi-person FPSB auctions. We will continue to assume that that each agent has only two possible strategies $\{b_1, b_2\}$.

Theorem 4.4.2. *In a multi-person symmetric FPSB auction with discrete bids and incomplete information, there exists a symmetric equilibrium as featured by equation (4.5), where the z value is given by*

$$z = \frac{\Pr(Y_{n-1} \le z)(b_2 - \frac{b_1}{n}) + \Pr(z \le Y_1)\frac{b_2}{n} + \sum_{i=1}^{n-2} \Pr(Y_{n-i-1} < z < Y_{n-i})\frac{b_2}{i+1}}{\Pr(Y_{n-1} \le z)\frac{n-1}{n} + \Pr(z \le Y_1)\frac{1}{n} + \sum_{i=1}^{n-2} \Pr(Y_{n-i-1} < z < Y_{n-i})\frac{1}{i+1}}. \tag{4.8}$$

Proof: Similar to the two-person cases, the utility functions, when agent 1 bids b_1, can be written as

$$u_1(v_1, z, b_1) = \Pr(Y_{n-1} \le z)\frac{1}{n}(v_1 - b_1).$$

When agent 1 bids b_2, there are two cases in which agent 1 wins the item. First, if all agents' valuations are less than z, agent 1 has probability $\Pr(Y_{n-1} \le z)$ to win the item. Second, if $Y_{n-i-1} < z < Y_{n-i}$, agent 1 has probability $\Pr(Y_{n-i-1} < z < Y_{n-i})$ to tie with $i+1$ agents. The utilities can be written as

$$\begin{aligned} u_1(v_1, z, b_2) &= \Pr(Y_{n-1} \le z)(v_1 - b_2) + \Pr(z \le Y_1)\frac{v_1 - b_2}{n} \\ &\quad + \sum_{i=1}^{n-2} \Pr(Y_{n-i-1} < z < Y_{n-i})\frac{v_1 - b_2}{i+1}. \end{aligned}$$

Agent 1 will bid b_1 if $u_1(v_1, z, b_1) \ge u_1(v_1, z, b_2)$; otherwise, it bids b_2. Subtracting $u_1(v_1, z, b_2)$ from $u_1(v_1, z, b_1)$, we obtain

$$\begin{aligned} u_1(v_1, z, b_1) - u_1(v_1, z, b_2) &= \Pr(Y_{n-1} \le z)\frac{1}{n}(v_1 - b_1) \\ &\quad - \Pr(Y_{n-1} \le z)(v_1 - b_2) - \Pr(z \le Y_1)\frac{v_1 - b_2}{n} \\ &\quad - \sum_{i=1}^{n-2} \Pr(Y_{n-i-1} < z < Y_{n-i})\frac{v_1 - b_2}{i+1}. \end{aligned}$$

Solving for v_1,

$$v_1 \le \frac{\Pr(Y_{n-1} \le z)(b_2 - \frac{b_1}{n}) + \Pr(z \le Y_1)\frac{b_2}{n} + \sum_{i=1}^{n-2} \Pr(Y_{n-i-1} < z < Y_{n-i})\frac{b_2}{i+1}}{\Pr(Y_{n-1} \le z)\frac{n-1}{n} + \Pr(z \le Y_1)\frac{1}{n} + \sum_{i=1}^{n-2} \Pr(Y_{n-i-1} < z < Y_{n-i})\frac{1}{i+1}}. \tag{4.9}$$

Combining equation (4.9) and equation (4.5), we obtain equation (4.8),

$$z = \frac{\Pr(Y_{n-1} \le z)(b_2 - \frac{b_1}{n}) + \Pr(z \le Y_1)\frac{b_2}{n} + \sum\limits_{i=1}^{n-2} \Pr(Y_{n-i-1} < z < Y_{n-i})\frac{b_2}{i+1}}{\Pr(Y_{n-1} \le z)\frac{n-1}{n} + \Pr(z \le Y_1)\frac{1}{n} + \sum\limits_{i=1}^{n-2} \Pr(Y_{n-i-1} < z < Y_{n-i})\frac{1}{i+1}}.$$

Thus, the z constructs an equilibrium strategy for an agent in a multi-person symmetric FPSB auction with discrete bids and incomplete information. $\diamond$

4.5 Conclusions

In this chapter, we discuss the existence and multiplicity of equilibria in FPSB auctions when bids are discrete. We point out that there are eleven different situations which lead to at most three equilibria and at least one equilibrium in 2-person FPSB auctions with complete information. In N-person FPSB auctions, we also show the existence of equilibrium and provide equilibrium solutions. We also discuss variations with sequential auctions. When the FPSB auctions have incomplete information, we provide equilibrium solutions for both two-person and multi-person FPSB auctions. In the incomplete information cases, we assume that agents have only two possible strategies. We expect that the computation procedure is similar to the cases when agents have more than two possible strategies. In those cases, nevertheless, we need to compute multiple different piecewise z values. And, the computational complexity will increase significantly.

Clearly, the random tie-breaking rule plays an important role in the refinement of the results. If we adopt different tie breaking rules, we may derive different results. This work uses a common tie-breaking rule, in which all tied bidders are equally likely to win. A possible alternative is to use Vickery auction to break ties, as Maskin and Riley [60] introduced to solve equilibria in FPSB auctions with continuous bids. It remains an area for future work to compare different tie-breaking rules in these models.

Chapter 5

Monte Carlo Approximation for Solving Sequential Auctions

In this chapter, we develop a more general approach to constructing trading agents based on game theory, and explore its computational limitations. We develop our technique in the context of a sequence of (possibly multi-unit) auctions with a small set of identified, risk-neutral participants, each of whom wants one unit of the item for which they have an independent, private value. We assume that our agent knows the distribution of the other agents' valuations, but not their actual values. This is meant to model common procurement scenarios, and may fit some markets on eBay in which it is apparently common for a small community of expert traders to recognize each other. In both situations, the relatively small number of significant opponents creates the opportunity to directly model one's competitors.

We cast the problem as an incomplete, imperfect information game. However, the straightforward expansion of a sequence of auctions creates a game that is intractable even for very small problems, and it is beyond the capability of current game solution software to solve for the Bayes-Nash equilibria. Thus, we construct a bidding policy through Monte Carlo sampling. In particular, we sample the opponents' valuations, assume they play perfectly, and solve the resulting imperfect information game. We accumulate the results of the sampling into a heuristic strategy for the incomplete information game.

The resulting strategy implicitly captures the belief updating associated with observing the opponents' bids in earlier auctions. Underlying this work is the assumption that information we gain about the other bidders can be used to improve play in later stages of the game. In particular, our observations of a bidder's actions in previous auctions should affect our belief about her valuation. For example, if we notice that Sue has placed bids at high values in previous auctions but not yet won anything, we are more likely to believe that Sue has a high valuation, which may influence how we should bid in future auctions.

The primary motivation of this line of work is to explore the potential benefits and the practical limitations of this approach. We find that the straightforward expansion of the imperfect information game cannot be solved directly by current game solvers (e.g., GAMBIT[1]). Thus, we develop methods to take advantage of the sequential structure that greatly reduces the space required to represent the game. Though this decomposition enables us to solve larger games, GAMBIT's ability to solve the decomposed games remains a bottleneck.

In Section 5.1 we formalize our model of the sequential auction scenario and set up the game theoretic analysis. Section 5.2 describes how we leverage the substructure to significantly decrease the amount of computation necessary to solve the game. In Section 5.3 we use Monte Carlo sampling to generate a heuristic bidding policy for our agent. Section 5.4 presents our empirical results, including comparisons between our heuristic policy and perfect play in markets that contain both single-unit and multi-unit auctions. Section 5.5 develops the relationship between our approach and the mathematics underlying sequential equilibria. We present related work in Section 5.6 and then conclude.

5.1 Model

Consider an agent, i, that has the task of purchasing one item from a sequence of auctions, K. Let c be the number of auctions, and k an individual auction. We refer to the collection of auctions as the *marketplace*. Individual auctions may offer multiple units and differ in the manner in which they form prices. The specification of the order and rules of the collection of auctions is the *market configuration*.

Let $q(k)$ be the number of units offered in auction k, and the total number of objects be

[1] The GAMBIT toolset is a software program and set of libraries that support the construction and analysis of finite extensive form and normal form multi-player games. See `http://www.hss.caltech.edu/gambit/Gambit.html`.

$q = \sum_k q(k)$. The auctions close in a fixed, known order, and in this model, all are treated as sealed bid auctions. The sealed bid assumption may not be as restrictive as it seems. In fact, the sniping strategy used by many bidders on eBay [89, 96] reduces the open-outcry auction to the equivalent of a sealed bid auction.

Let J denote the other bidders in the market, and $A = J \cup i$. The total number of bidders, including i, is $n = |J| + 1$. In a particular auction, a subset, $\mathcal{A} \subseteq A$, of the agents will place bids. Let the bid of bidder j in auction k be denoted b_j^k.

Naturally, the rules of the auctions will affect the bidders' choices of actions. A multi-unit auction must have a policy for setting prices.[2] In this study, we consider only two such policies. The *Mth-price* policy sets the price paid by all winners to the value of the lowest winning bid (this is the policy used in eBay's Dutch Auction format). Under the *pay-your-bid* policy, each winner pays the price she offered. Pay-your-bid is the policy used on Yahoo's multi-unit auctions. In the case of a single unit for sale, the two policies are equivalent.

Given a sequence of sealed-bid auctions, the agent must select a bid to place in each auction. Let B^k be the set of bid choices that are acceptable in auction k. Typically, we assume that B^k is the set of integers in some range and is identical across all of the auctions. However, the techniques we develop admit different bid choices in each auction. The number of bid choices is $m = |B^k|$. We assume that ties are broken randomly.

Our agent has a value $v_i(k)$ for an item in k, and bidder $j \in J$, has valuation $v_j(k)$. In this study, we assume that the items available in K are identical and that all participants are interested in only a single unit. We anticipate that the techniques we develop in this chapter can be extended to auctions of heterogeneous items if an agent's valuations for the items are correlated, that is, if learning about an agent's valuation of one item helps predict its valuation of another item.

Agent i does not know bidder j's true value for the items, but knows that it is drawn from a distribution, D_j. In this model, we assume that valuations are independent and private, but we do not make any particular assumptions about the functional form of the distributions, nor do we assume that the distributions are identical for all of the bidders. We will make various assumptions about whether the bidders in J know each other's valuations or agent i's valuation.

We assume that each participant is present for the first auction, and continues to participate in each auction until either she wins or the sequence ends. Thus, a buyer that does not win in auction k will participate in auction $k + 1$. We assume that the auctioneer makes public a list of all of the

[2] See [111] for a survey of some pricing policies.

bids once the auction is complete. This is consistent, for instance, with eBay's policy. Let h_j^k be the sequence of bids that agent j placed in the auctions up to, but not including, k. That is, $h_j^k = \{b_j^1, \dots, b_j^{k-1}\}$. We call h_j^k bidder j's *history* up to auction k. The history of all J bidders leading to auction k is denoted H_J^k.

5.1.1 Sequential Game Representation

We model the sequential auction scenario as an extensive form game, $\Gamma(\mathcal{A}, V_{\mathcal{A}}, B^K, K)$, where $\mathcal{A} = J \cup i$ and B^K denotes the bid choices for all of the auctions. A *subgame* has the same structure, except that part of the game has already been played. For example, the subgame that results when bidder j wins the first item is $\Gamma(\mathcal{A}', V_{\mathcal{A}'}, B^{K'}, K')$ where $\mathcal{A}' = i \cup J \setminus j$ and $K' = K \setminus \{1\}$.

It is also useful to identify the game structure of individual auctions. Denote a *component* auction game $\gamma(\mathcal{A}, V_{\mathcal{A}}^k, B^k)$, in which agents $\mathcal{A}$, with valuations $V_{\mathcal{A}}^k$ for the items in auction k, choose bids from the domain B^k. Note that a game (or subgame) is a sequence of component games. In game theoretic terms, γ is the game in which $\mathcal{A}$ is the set of players, B^k are the actions, and the payoff is $v_j(k) - b_j^k$ for the bidder with the highest bid, and zero for everyone else. Because the auction is sealed bid, all of the bidders' actions are simultaneous, and the game involves imperfect information.

A simple example with three agents, two items, and two bid levels is shown in Figure 5.1. The circles are labeled with the ID of the agent, and the arcs with the bid value ($\{1, 2\}$). The game consists of two stages, the first of which corresponds to the first auction involving all three agents. The second stage involves the two agents who did not win the first item, and for conciseness, we have substituted labeled triangles for subgames on the leaves of the first auction. There are fifteen subgames, labeled $\gamma_1 \dots \gamma_{15}$, but only three possible unique structures, labeled A, B, and C.

Dotted lines connect decision nodes in the same information state. The small squares at the leaves of the subgames represent terminal states that would be labeled with the payoffs to the agents. The actual value of the payoffs would depend upon each agent's actual value for the item, the path taken, and the auction's policy for setting prices. The diamonds denote the random move by nature to break ties among the bids (with the probabilities indicated in parenthesis). This type of move by nature can be handled relatively easily because it does not introduce any asymmetric information. Moreover, it is amenable to the decompositions we introduce in the next section.

It is obvious from Figure 5.1 that a particular component game, γ, can appear many times in the overall game Γ. Each second level component game appears on five different paths of the top

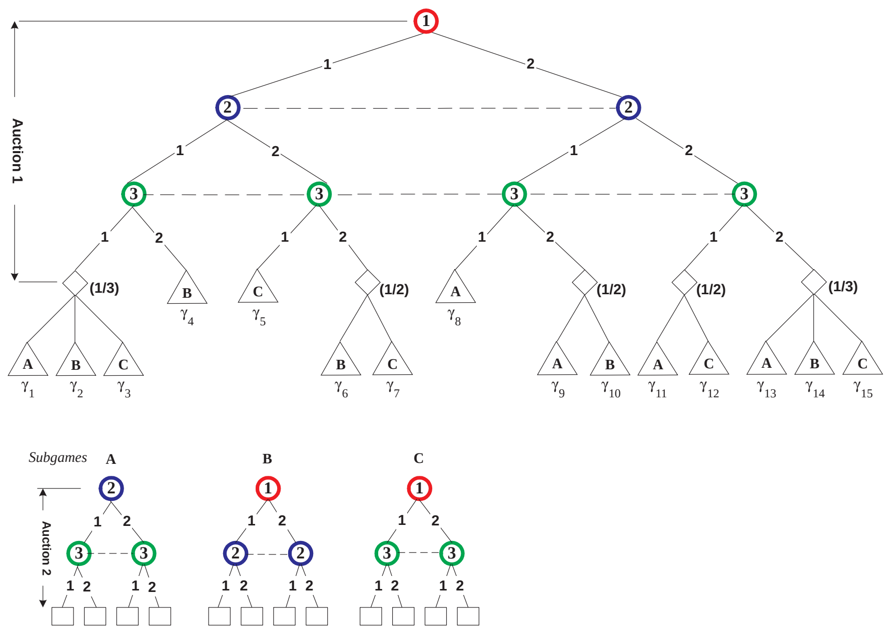

Figure 5.1: A sequence of two sealed-bid auctions with three agents, one item for sale in each auction, and two bid levels.

level game. When necessary, we will distinguish a component game using its history as a subscript: $\gamma_{H_J^k}$. The history information is sufficient to uniquely identify each component game instance.

In addition to the imperfect information generated by the sealed bids, the agent also faces incomplete information because it does not know the other bidders' true values, and therefore does not know the other bidders' payoffs. Harsanyi (1967-68) demonstrated that incomplete information games can be modeled by introducing an unobservable move by nature at the beginning of the game which establishes the unknown values. This approach transforms the incomplete information game into a game with imperfect information.

Unfortunately, the move-by-nature approach is computationally problematic. The number of possible moves available to nature is m^n, where m is the size of the domain of $v_j(k)$, and n is the number of agents. Our model permits a continuous range for valuation functions, so the number of choices is not enumerable. In some special cases, analytic solutions can be found to auction games with continuous types [24]. However this analysis is complex and typically requires restrictive assumptions about the distributions of values. Moreover, whether valuations are drawn from discrete or continuous domains, each different market configuration requires a separate analysis.

For these reasons, we investigate the use of Monte Carlo sampling to generate heuristic bidding policies for the incomplete information game. Our approach to the problem can be summarized as follows:

1. Create a sample complete-information game by drawing a set of valuations for other bidders.
2. Solve for a Nash equilibrium of the sample game.
3. Update the agent's bidding policy.

The first step is straightforward Monte Carlo sampling. The second and third steps are the subject of the next two sections.

5.2 Leveraging Substructure in the Complete Information Game

We built our agent on top of the GAMBIT Toolset. Although GAMBIT includes algorithms that can solve multi-player games with imperfect information, it cannot solve the straightforward expansion of even very small instances of the complete-information, sequential auction game in a reasonable amount of time.

To see why, consider the size of the extensive form of a complete information sequential auction game with ties broken randomly. The assumption that bidders want only one item means that the winners of a particular auction will not participate in future auctions. Thus, auction $k+1$ has $q(k)$ fewer participants than auction k. In general, the number of agents participating in component game k is $z(k) = n - \sum_{x=1}^{k-1} q(x)$. The number of nodes in the extensive form representation of this game with c auctions is

$$\frac{m^n - 1}{m-1} + \sum_{k=2}^{c} \left[\frac{m^{z(k)} - 1}{m-1} \times \prod_{j=1}^{k-1} \left(m^{z(j)} + \mathrm{EXT}\left[z(j), m, q(j)\right] \right) \right].$$

The core of the equation captures the number of nodes in the tree without tie breaking, and the EXT term represents the number of additional terminal nodes added to each component game due to tie breaking. The EXT term expands as

$$\begin{aligned}
&\mathrm{EXT}\left[z(j), m, q(j)\right] \\
&\quad = \sum_{v=1}^{m} \sum_{i=q(j)+1}^{z(j)} \binom{z(j)}{i} (v-1)^{z(j)-i} \left[\binom{i}{q(j)} - 1 \right] \\
&\quad\quad + \sum_{v=1}^{m} \sum_{i=2}^{z(j)-1} \binom{z(j)}{i} \sum_{h=L(j,i)}^{H(j,i)} \binom{z(j)-i}{h} (m-v)^h (v-1)^{z(j)-i-h} \left[\binom{i}{q(j)-h} - 1 \right],
\end{aligned}$$

where

$$H(j,i) = \min(q(j)-1, z(j)-i), \text{and}$$
$$L(j,i) = \max(q(j)-i+1, 1).$$

A five agent, four item sequential auction with five bid choices and random tie breaking has 4.5 billion decision nodes and is unsolvable with GAMBIT on current workstations. However, as Figure 5.1 suggests, there is structure in the problem that we can leverage to improve our representation of the game.

The computational aspects of game theory have been studied by economists and computer scientists in the past few years [44, 45, 46, 66, 105]. A very promising thread of work is focused on representations of games that capture their inherent structure and facilitate solution computation. Koller and Pfeffer's GALA language (1997) can be used to represent games in *sequence* form, and the authors have developed solution techniques for two-player, zero-sum games represented in this format. The success of GALA is based on the intuition that significant computational savings can be achieved by taking advantage of a game's substructure. This intuition holds for the sequential auction model, and we have employed it to improve upon GAMBIT's default approach.

The default representation of this game in GAMBIT is to expand each of the leaves with an appropriate subgame. Given that the bidders have complete information, all subgames with the same players remaining have the same solution(s). Thus, a single-unit, sealed-bid (component) auction with n agents has at most n *unique* subgames—one for each possible set of non-winners. The three component games—A, B, and C—are illustrated in Figure 5.1.

Our agent's approach is to create all possible component games and solve them using GAMBIT's C++ libraries. The process is essentially dynamic programming, and equivalent to standard backward induction with caching. The expected payoffs from the solution to a component game γ involving bidders $\mathcal{J}$ are used as the payoffs for the respective agents on the leaves of any component games in Γ which immediately precede γ. The agent solves all possible smallest component games (i.e., where $k = c$), and recursively constructs higher-order subgames until it solves the root game (i.e., $k = 1$).

The number of decision nodes required to express a game in its component form is

$$\sum_{k=1}^{c} \binom{n}{z(k)} \frac{m^{z(k)} - 1}{m - 1}.$$

The component form representation is exponential in the number of agents and the number of bidding choices. However, the total number of nodes required to express the game is exponentially less than in the full expansion. For example, the five agent, four item sequential single-unit auctions with five bid choices and random tie-breaking requires only 1931 nodes to encode in its component form, compared to the 4.5 billion required for the naive expansion.

It should be noted that the solutions that we are using in the above analysis are Nash equilibria found by GAMBIT for each particular subgame. These solutions may involve either pure or mixed strategies. It is well known [82], that at least one mixed strategy equilibrium always exists, however it is also often true that more than one Nash equilibria exist. In this study, we simply take the first equilibria found by GAMBIT, and leave the question of how, and even whether, to incorporate multiple equilibria to future research. We recognize that our results may be influenced by the order in which GAMBIT finds solutions, but also consider it a concern inherent in using off-the-shelf solution technology.

It should also be noted that the procedure described above is consistent with the definition of *subgame perfect equilibrium (SPE)*, a well-known specialization of Nash equilibria. A profile of strategies is subgame perfect if it entails a Nash equilibrium in every subgame of the overall game [94]. All subgame perfect equilibria are Nash, but the reverse is not necessarily true.

While the decomposition provides an exponential improvement in the number of nodes needed to represent (and hence solve) the game, the computational cost of finding equilibria for the component games remains a severely limiting factor. Indeed, though the number of bid choices is the base, not the exponent, of the complexity of the extensive form game, we will see in Section 5.4 that GAMBIT is unable to solve subgames if we increase the number of bid choices beyond a small number.

5.3 Monte Carlo Approximation

In order to participate in this environment, the agent must construct a *policy*, Π, that specifies what action it should take in any state of the game that it might reach. There are many conceivable policies available to our agent.

One simple strategy is to compute the equilibrium strategy in each component game, and to bid accordingly. For example, the equilibrium strategy of a single first-price, sealed-bid auction in which the other bidders' valuations are drawn uniformly from $[0, 1]$ is to bid $b_i^k = (1 - 1/n)v_i(k)$, where n is the number of bidders [62]. We define Π_{myopic} to be the strategy in which the agent bids according to the equilibrium of each individual sealed-bid auction. Thus, the strategy has one element for each potential game size, $\Pi_{\text{myopic}} = \{\pi_z\}$ where z is the size, in number of bidders, of the component game.

In a sequence of sealed-bid, single-unit auctions, a Bayes-Nash equilibrium strategy is for a bidder to bid the expected price of the $(q+1)$st valuation under the assumption that her bid is among the top q (see [106] for details). We denote this policy $\Pi_{(\text{q+1})\text{st}}$ and use it as a benchmark in our empirical evaluation.

If the distributions from which the bidders draw values are not identical, then it would behoove our agent to have a policy that accounted for which other bidders were in the subgame. Thus, $\Pi_{\text{not-id}} = \{\pi_{\mathcal{J}\subseteq J}\}$. That is, the actions in the policy depend upon which subset, $\mathcal{J}$, of agents remain.

All three policies mentioned thus far are memoryless; they ignore the bids the remaining opponents made in previous auctions. On the other extreme is a policy that uses all possible history information. $\Pi_{\text{history}} = \{\pi_{\mathcal{J},H^k_J}\}$ encodes the entire tree because the decision at each decision node is a function of the entire history.

The policy that our agent learned in this study is $\Pi_{\text{agg-hist}} = \{\pi_{\mathcal{J},H^k_{\mathcal{J}}}\}$ where $H^k_{\mathcal{J}} = \{h^k_{j\in\mathcal{J}}\}$, the histories of all other agents still in the game. This differs from Π_{history} in that policies are classified by the histories of only those bidders that remain active ($\mathcal{J}$), rather than by the previous actions of all bidders in J. It is based on the assumption that bidders who are no longer active in the sequential auction (because they have won an item) are irrelevant. Therefore, all component games that have the same opponents and identical previous actions by those agents, are aggregated into a class of component games, $\gamma_{\mathcal{J},H^k_{\mathcal{J}}}$.

In the example in Figure 5.1, suppose Player 1 is our agent. All paths that lead to subgame A can be ignored because our agent won the item in the first auction. Of the remaining subgames, the set $\{\gamma_2, \gamma_4, \gamma_{10}\}$ have identical histories—bidder 2 bid \$1 in all of them. Similarly, the sets $\{\gamma_6, \gamma_{14}\}$, $\{\gamma_3, \gamma_5, \gamma_{12}\}$, and $\{\gamma_7, \gamma_{15}\}$ can be formed by their common histories.

The agent constructs the policy by sampling the distributions of the other bidders and solving the resulting complete information game. Let L be the collection of sample games constructed, and l a single instance. Denote the solution returned by GAMBIT to instance l as Ω^l. Ω^l is a profile of (possibly mixed) strategies—one for each player—that constitute an equilibrium for this game instance. Let, Ω^l_i specify the policy for agent i, and $\omega^l_i(\gamma)$ is the policy for subgame γ. Note that some decision nodes may not be reachable if the actions that lead to them are played with zero probability. To simplify the notation, we include these unreachable nodes in the following even though they have no effect on the solution.

To compute the policy $\pi_{\mathcal{J},H^k_{\mathcal{J}}}$ for a decision in game $\gamma_{\mathcal{J},H^k_{\mathcal{J}}}$ we take the weighted sum of the equilibrium solutions across all sample games. Let

$$w(b^k_i|\pi_{\mathcal{J},H^k_{\mathcal{J}}}) = \sum_{l \in L} \sum_{\gamma \in \gamma_{H^k_{\mathcal{J}}}} \Pr(\gamma|\Omega^l) \Pr(b^k_i|\omega^l_i(\gamma)) \tag{5.1}$$

be the weight assigned to action b^k_i in the class of games identified by $\gamma_{H^k_{\mathcal{J}}}$. Here, $\Pr(\gamma|\Omega^l)$ is the probability that the game would reach subgame γ given that everyone is playing Ω^l (i.e., the product of the probabilities in the mixed strategies on the path leading to γ, and $\Pr(b^k_i|\omega^l_i(\gamma))$ is the probability associated with bid b^k_i in solution $\omega^l_i(\gamma)$.

Zhu and Wurman examine a version of the update function with a bias towards actions that generate a higher utility for our agent [115]. The inclusion of utility in the equation biases the agent toward maximizing its expected utility—a useful heuristic, perhaps, but one that is not necessarily consistent with equilibrium behavior. In this work, we compare the effect of using the biased update function rather than the unbiased equation (5.1). The biased updated function has the form:

$$w(b^k_i|\pi_{\mathcal{J},H^k_{\mathcal{J}}}) = \sum_{l \in L} \sum_{\gamma \in \gamma_{H^k_{\mathcal{J}}}} \Pr(\gamma|\Omega^l) u_i(\gamma,\Omega^l) \Pr(b^k_i|\omega^l_i(\gamma)), \tag{5.2}$$

where $u_i(\gamma,\Omega^l)$ is our agent's expected utility of the subgame rooted at γ.

Finally, we normalize the computed weights to derive the probabilities,

$$Pr(b^k_i|\pi_{\mathcal{J},H^k_{\mathcal{J}}}) = \frac{w(b^k_i|\pi_{\mathcal{J},H^k_{\mathcal{J}}})}{\sum_{b \in B^k} w(b|\pi_{\mathcal{J},H^k_{\mathcal{J}}})}. \tag{5.3}$$

The result of this process is a policy that specifies a (possibly mixed) strategy for each unique class of component games. We refer to a policy constructed in this manner as a *Monte Carlo Approximation* (MCA) policy.

5.4 Empirical Results

To evaluate the efficacy of the approach, we simulated several market configurations in which we varied the functional form of the valuation distributions, the form of the update equation, and the strategies of the other bidders. Each of these experimental variables are described in more detail below. The experimental design is similar to the previous work by Zhu and Wurman [115]. However, in the results reported herein, we have added random tie-breaking rule and multi-unit auctions.

- **Market Configuration:** The market configuration includes the number of agents, the domain of the bid messages, and the number and types of auctions. We used the following configurations:
  - $\{$5,5,s-s-s$\}$ contains five agents, five bid levels, and a sequence of three single-item auctions.
  - $\{$5,5,s-2Mth$\}$ contains five agents, five bid levels, and an auction sequence in which a single-unit auction is followed by a Mth-price auction for two units.
  - $\{$5,5,s-2PYB$\}$ contains five agents, five bid levels, and an auction sequence in which a single-unit auction is followed by a two-unit auction in which the winners pay their bid values.
  - $\{$5,4,s-s-s-s$\}$ contains five agents, four bid levels, and a sequence of four single-item auctions.
- **Valuation Distribution:** we used three types of distributions: uniform, left-skewed Beta, right-skewed Beta. With the exception of $\{$5,4,s-s-s-s$\}$, the valuations of the other agents were drawn from [1, 6], while our agent's valuation is always fixed at 3.5. In the left-skewed distribution, our agent is likely to have a valuation significantly above average, while in the right-skewed distribution it will be significantly below average. In experiments with $\{$5,4,s-s-s-s$\}$, the valuations of the other agents drawn from [1,5] while our agent's valuation is fixed at 3; this combination was chosen to draw comparisons with an earlier work by Zhu and Wurman [115].
- **Update Equation:** we examined the difference between using equation (5.1) and using equation (5.2), which biases the policy aggregation by the agent's expected utility.
- **Bidder Strategies:** we studied the effects of various combinations of bidder strategies.
  - *All SPE:* as a benchmark scenario, we assume that all agents have complete information for a test case and all of them play the subgame perfect equilibrium computed using our structural decomposition technique with the GAMBIT engine.
  - *MCA/n-SPE:* we assume the other agents had complete information, while our agent has incomplete information. Our agent implements the strategy learned from the Monte Carlo policy construction, while the other agents implement their SPE strategies. Since

our agent is not playing perfectly, there is no guarantee that the other agents' SPE strategies are equilibrium responses to our imperfect play.[3] To generate the MCA strategy, the agent trained with 200 samples.

– *All MCA:* In this scenario all agents construct and play strategies generated with Monte Carlo policy construction. Note that for these simulations, each opponent must be retrained with each new draw of its valuation.

– $(q + 1)$-*Equilibrium:* Another benchmark for the sequence of single-unit auctions, in the $(q+1)$-equilibrium strategy all agents play the sequential auction equilibrium strategy [106]. Each agent bids the expected price of the $(q + 1)$st valuation under the assumption that its bid is among the top q.

In the experiments, we measure the utility for our agent (computed as the difference between its value and the price it pays if it wins), the social welfare (the aggregate value of all of the winning agents), and the revenue achieved by the seller. The experiments were run on a Beowulf cluster of eight Linux computers.

In some cases, our agent may find that the game has progressed down a path for which it learned no policy. In such cases, our agent picks the most similar subgame for which it does have a policy. The similarity measure favors subgames with the same bidding pattern, but possibly different agents, over subgames with the same agent but different bidding patterns.

Figure 5.2 shows our agent's utility on thirty randomly selected problem instances from the $\{$5,4,s-s-s-s$\}$ market scenario with other agents' valuations drawn from the uniform distribution. For each problem instance, the four strategy combinations were tested, and update equation (5.2) is used. The performance of the Monte Carlo strategy is quite close to that of the subgame perfect equilibrium both when the other agents play perfectly and when they construct their own Monte Carlo strategies. From this result we conclude that the approximation technique generates policies that perform quite well in this environment.

The $(q + 1)$-equilibrium strategy is included in Figure 5.2, though it is important to note that it represents a slightly different game than the other three. Agents must be allowed to place real-valued bids in the $(q + 1)$-equilibrium strategy, while in the other three we are restricting bids to integer values. This distinction explains, for instance, why our agent achieves zero utility in Figure 5.2 under the $(q + 1)$-equilibrium strategy when it has the lowest value among the five

[3] In theory, it would be possible to determine the opponents' best responses to our heuristic strategy by marginalizing our agent and computing a reduced game in which the other agents' payoffs are impacted by our fixed behavior.

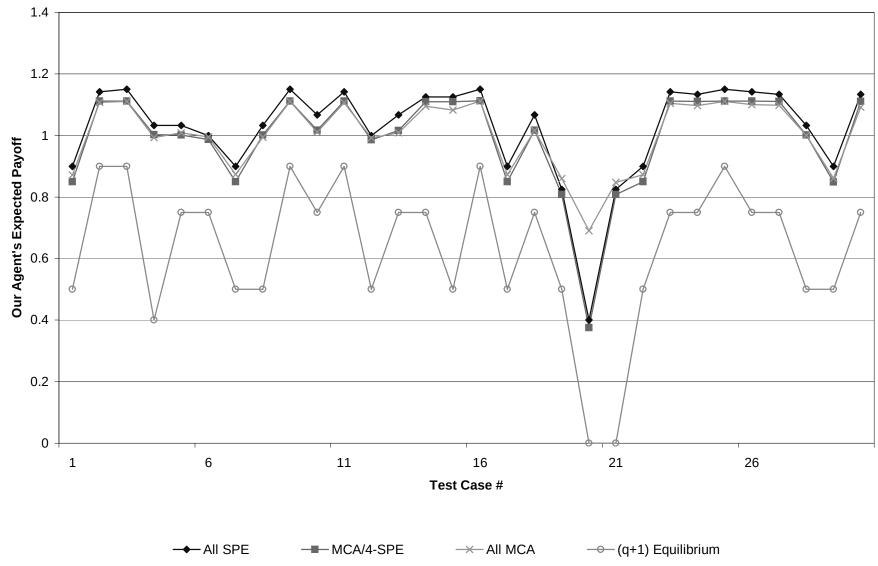

Figure 5.2: Our agent's expected payoff in the {5,4,s-s-s-s} market scenario with the other agents' valuations drawn from a uniform distribution and equation (5.2) is used to update policies.

agents. When bid values are restricted, it is more likely that our agent will end up in a tie and therefore achieve a positive surplus with some probability. Nevertheless, the pattern of the payoffs for the $(q+1)$-equilibrium strategy is quite similar to our empirical results.

One aspect that we want to examine is the effect of the utility term in equation (5.2). Figure 5.3 shows our agent's expected utility on the same 30 test cases when trained with the same training data and equation (5.1). Although Figures 5.2 and 5.3 look nearly identical, close inspection shows that equation (5.2) performs slightly better than equation (5.1), in the sense that it more closely approximates the subgame perfect outcomes. For this reason, we continue to use equation (5.2) in the rest of the empirical tests.

Figures 5.4 and 5.5 show similar correspondence between the strategies when the other agents' valuations are drawn from right-skewed and left-skewed Beta distributions, respectively. Notice that the in the left-skewed distribution our agent achieves higher payoffs, while in the right-skewed case our agent receives lower payoffs. This result is expected given that the expected average valuation will be lower when the opponents are drawn from a left-skewed distribution, and higher

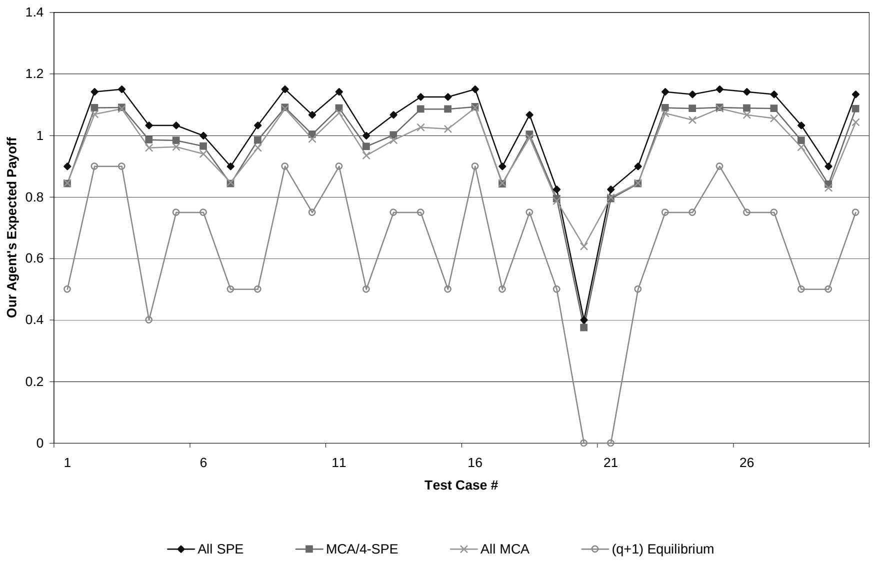

Figure 5.3: Our agent's expected payoff in the {5,4,s-s-s-s} market scenario with the other agents' valuations drawn from a uniform distribution and equation (5.1) is used to update policies.

when drawn from a right-skewed distribution.

The next set of experiments involved five-agent, three-item scenarios. We compared two multi-unit auction scenarios, {5,5,s-2Mth} and {5,5,s-2PYB}, against a sequence of three single unit auctions, {5,5,s-s-s}, over the same thirty uniform-distribution sample instances tested above. Figures 5.6 and 5.7 show how closely the performance of the MCA strategy tracks that of the subgame perfect strategy for {5,5,s-2Mth} and {5,5,s-2PYB}, respectively. Figure 5.8 contrasts our agent's payoff for the three scenarios. The results from {5,5,s-2Mth} and {5,5,s-2PYB} are nearly identical (and may appear to be a single line), while significant variation exists in results from {5,5,s-s-s}. Notice that our agent performed significantly better in both {5,5,s-2Mth} and {5,5,s-2PYB} than in {5,5,s-s-s}. It is clear that, overall, the agents are bidding lower in the multi-unit scenarios, and our agent is playing a mixed strategy that is more successful. However, it remains to be seen whether there is a game theoretic explanation for this outcome, or whether it is a byproduct of our technique or the manner in which GAMBIT returns solutions.

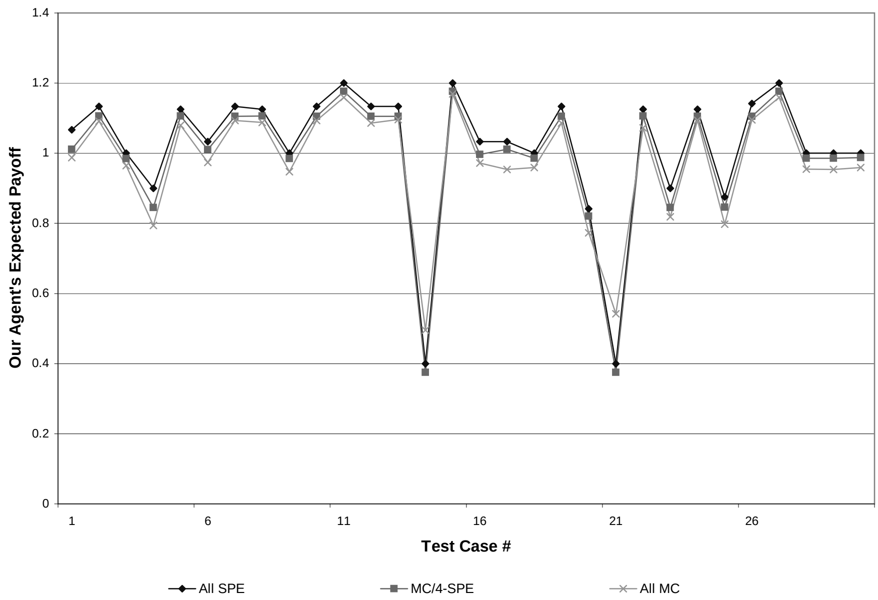

Figure 5.4: Our agent's expected payoff in the {5,4,s-s-s-s} market scenario with the other agents' valuations drawn from a right-skewed Beta distribution.

Figure 5.9 shows the social welfare achieved in all three scenarios. The welfare achieved in scenario {5,5,s-s-s} is slightly better than the two multi-unit cases, whose graphs are again nearly coincident. This is consistent with the observation that the agents are behaving more collaboratively in the multi-unit auction by bidding lower and letting the tie-breaking determine the winner. When the agent with the highest value allows the allocation to be determined by tie-breaking rather than by placing a better bid, it is more likely that a less than optimal allocation will result.

Figure 5.10 shows the effect of the different auction scenarios on the sellers' revenue. Again, because buyers are acting more competitively in the single-unit auctions, the sellers achieve greater revenue than in the multi-unit auction scenarios.

5.5 Convergence of MCA Policies

A perfect Bayesian equilibrium is defined in terms of beliefs at decision points in the game, and requires that an equilibrium policy be consistent with those beliefs. In this section, we show that

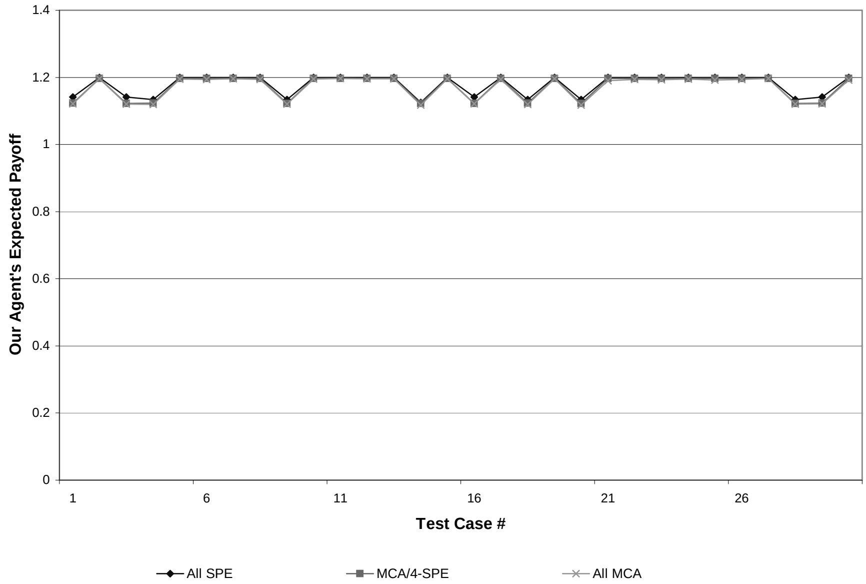

Figure 5.5: Our agent's expected payoff in the {5,4,s-s-s-s} market scenario with the other agents' valuations drawn from a left-skewed Beta distribution.

the MCA policy at a node implicitly captures the agent's beliefs about which opponent valuations would explain the fact that the agent arrived at a particular decision point in the game tree.

Theorem 5.5.1. *MCA converges to the average policy of perfect Bayesian equilibrium.*

Proof: Let Ω^V be a perfect Bayesian equilibrium profile of the game when agents have valuation profile V. Let Φ be our agent's belief function, and $\Phi(V)$ be our agent's belief that the other agents have valuation profile V. Let ϑ be an element of V. We have $\Pr(\Omega^\vartheta) = \Phi(\vartheta)$ and $\Pr(\Omega^V) = \Phi(V)$ for one specialization ϑ of V. Similarly, for a component game, we have $\Pr(\omega_i^l(\gamma)) = \Pr(\gamma)$. Let $\Pr(H_J^k|\Omega^V)$ be the probability that the policies selected by Ω^V follow history H_J^k. Given history H_J^k, the probability that the other agents have profile V, is given by

$$\Pr(V|H_J^k) = \frac{\Pr(H_J^k|\Omega^V)\Phi(V)}{\int_\vartheta \Pr(H_J^k|\Omega^\vartheta)\Phi(\vartheta)d\vartheta}.$$

A perfect Bayesian equilibrium will define a policy for a subgame that is consistent with the beliefs. Here we simply let the policy be the *average policy*, that is, the policy constructed by

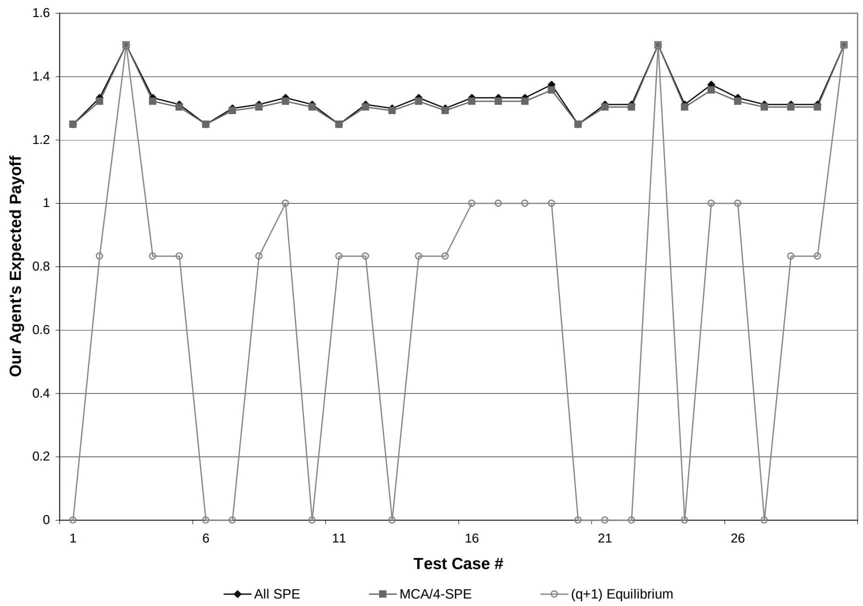

Figure 5.6: Our agent's expected payoff in the $\{$5,5,s-2Mth$\}$ scenario with the other agents' valuations drawn from a uniform distribution.

taking an average over all action profiles, weighted by the likelihood of seeing V given that we have reached the subgame. In other words, the probability that our agent plays b_i^k in subgame $\gamma_{H_J^k}$ is

$$Pr(b_i^k|H_J^k) = \int_\vartheta \Pr(b_i^k|\vartheta)\Pr(\vartheta|H_J^k)d\vartheta.$$

substituting $\Pr(\vartheta|H_J^k)$ with $\Pr(V|H_J^k)$ into the above equation, we have

$$\begin{aligned} Pr(b_i^k|H_J^k) &= \int_\vartheta \Pr(b_i^k|\vartheta)\frac{\Pr(H_J^k|\Omega^\vartheta)\Phi(\vartheta)}{\int_\vartheta \Pr(H_J^k|\Omega^\vartheta)\Phi(\vartheta)d\vartheta}d\vartheta \\ &= \frac{\int_\vartheta \Pr(b_i^k|\vartheta)\Pr(H_J^k|\Omega^\vartheta)\Phi(\vartheta)d\vartheta}{\int_\vartheta \Pr(H_J^k|\Omega^\vartheta)\Phi(\vartheta)d\vartheta} \\ &= \frac{\int_\vartheta \Pr(b_i^k|\vartheta)\Pr(H_J^k|\Omega^\vartheta)\Phi(\vartheta)d\vartheta}{\Pr(H_J^k)}. \end{aligned} \tag{5.4}$$

Let us consider MCA strategies. In MCA, we compute the policy by taking the weighted sum of the subgame perfect equilibrium solutions across all sample games. Let l be a sample in the collection of samples L. We know that each l maps to a V. So, $\Pr(l) = \Pr(V)$. $\omega_i^l(\gamma)$ is the

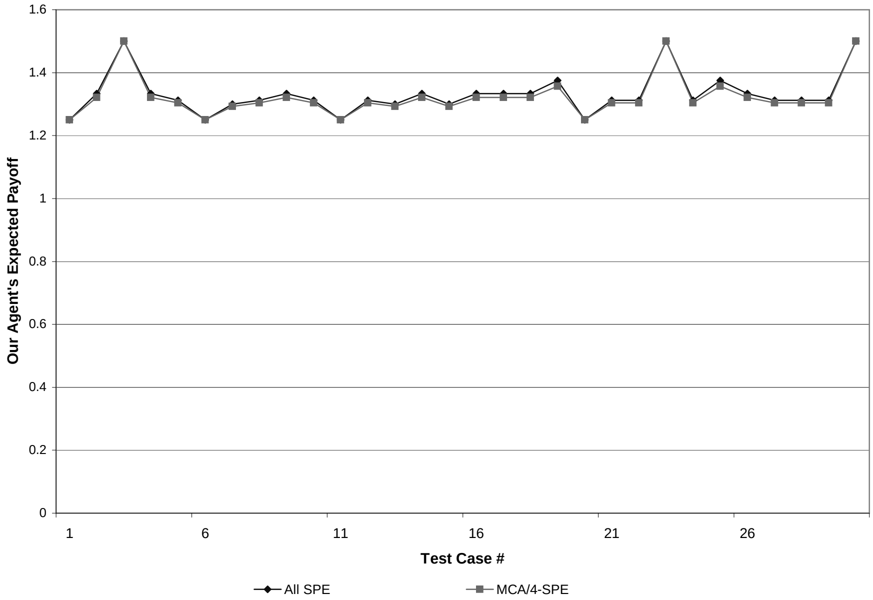

Figure 5.7: Our agent's expected payoff in the $\{$5,5,s-2PYB$\}$ scenario with the other agents' valuation are drawn from a uniform distribution.

solution to γ. So, $\Pr(b_i^k|\omega_i^l(\gamma)) = \Pr(b_i^k|\gamma)$. Like equation (5.1), we have

$$w(b_i^k|H_J^k) = \sum_{l\in L}\sum_{\gamma\in\gamma_{H_J^k}} \Pr(\gamma|\Omega^l)\Pr(b_i^k|\gamma). \tag{5.5}$$

Using the multiplication and Bayesian rules, we have

$$\begin{aligned} w(b_i^k|H_J^k) &= \sum_{l\in L}\sum_{\gamma\in\gamma_{H_J^k}} \Pr(b_i^k|\gamma)\Pr(\gamma|\Omega^l) \\ &= \sum_{l\in L}\sum_{\gamma\in\gamma_{H_J^k}} \Pr(b_i^k|H_J^k)\Pr(H_J^k|\gamma)\Pr(\gamma|\Omega^l) \\ &= \sum_{l\in L}\Pr(b_i^k|H_J^k)\Pr(H_J^k|\Omega^l) \\ &= \Pr(b_i^k|H_J^k)\sum_{l\in L}\Pr(H_J^k|\Omega^l). \end{aligned} \tag{5.6}$$

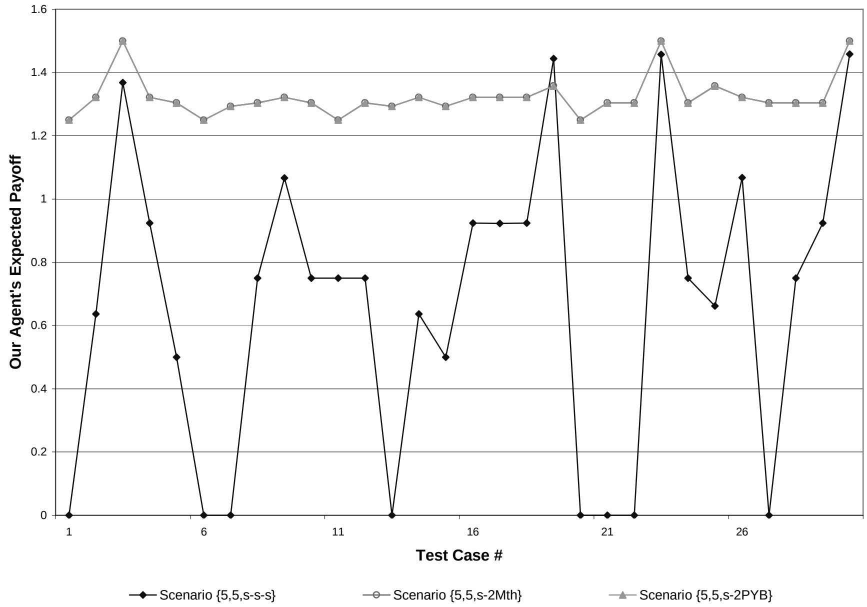

Figure 5.8: Comparison of our agent's expected payoff among different types of auctions by using MCA strategy while the other agents' valuation are drawn from a uniform distribution.

In MCA, $\Pr(b_i^k|H_J^k)$ is an equation described in numerical format as follows

$$\begin{aligned}\Pr(b_i^k|H_J^k)\sum_{l\in L}\Pr(H_J^k|\Omega^l) &= \frac{\Pr(H_J^k|b_i^k)\Pr(b_i^k)}{\Pr(H_J^k)}\sum_{l\in L}\Pr(H_J^k|\Omega^l)\\ &= \frac{\sum\limits_{l\in L}\Pr(H_J^k|\Omega^l)\Pr(\Omega^l|b_i^k)\Pr(b_i^k)}{\Pr(H_J^k)}\sum_{l\in L}\Pr(H_J^k|\Omega^l)\\ &= \frac{\sum\limits_{l\in L}\Pr(H_J^k|\Omega^l)\Pr(b_i^k|\Omega^l)\Pr(\Omega^l)}{\Pr(H_J^k)}\sum_{l\in L}\Pr(H_J^k|\Omega^l). \qquad (5.7)\end{aligned}$$

Replace Ω^l with a value sample ϑ, combine equations (5.5),(5.6), (5.7), and let the number of samples go to infinity, we have $\sum\limits_{l\in L}\Pr(H_J^k|\Omega^l) = 1$ and

$$w(b_i^k|H_J^k) = Pr(b_i^k|H_J^k) = \frac{\int_\vartheta \Pr(b_i^k|\vartheta)\Pr(H_J^k|\Omega^\vartheta)\Phi(\vartheta)d\vartheta}{\Pr(H_J^k)}. \qquad (5.8)$$

Equation (5.8) is the same as the average policy of perfect Bayesian equilibrium in equation (5.4).

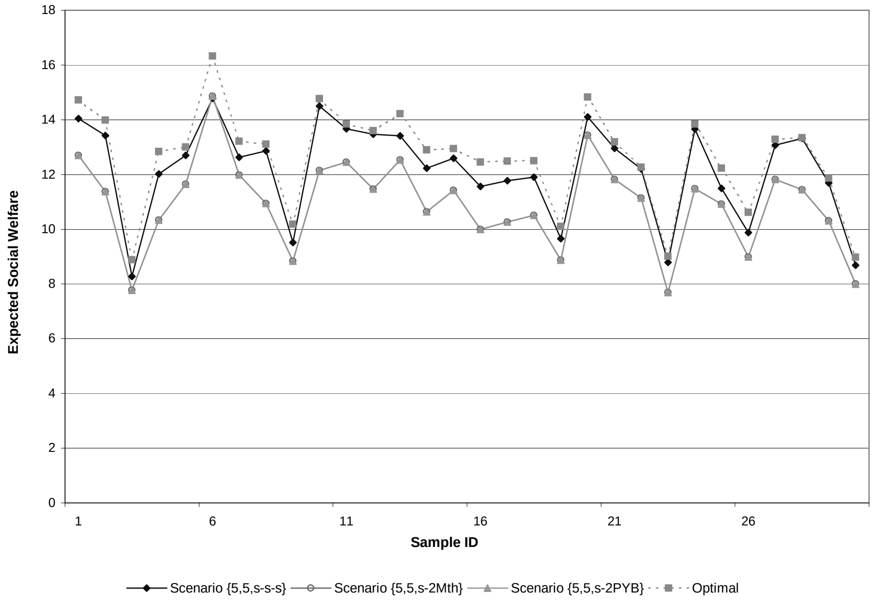

Figure 5.9: Comparison of the expected social welfare among different auction scenarios when our agent plays its MCA strategy and the other agents' valuation are drawn from a uniform distribution.

Thus, we conclude that MCA converges to the average policy of perfect Bayesian equilibrium when the number of samples goes to infinity.◇

5.6 Related Work

This work continues the study begun by Zhu and Wurman [115], which studied single unit sequential auctions with deterministic tie-breaking. In this work, we admit multi-unit auctions, random tie-breaking rules, and slightly larger problem sizes. Moreover, we connect the MCA approach directly to belief updating and sequential equilibria.

Our main focus is to study the feasibility of using game theory as a solution tool in a computational agent adaptable to various electronic market configurations. The copious research on auctions and game theory provides a backdrop for our effort. See Klemperer [42] for a broad review of auction literature, including a discussion of sequential auctions for homogeneous objects. Weber [106] shows that the equilibrium strategies for the bidders when the objects are sold in

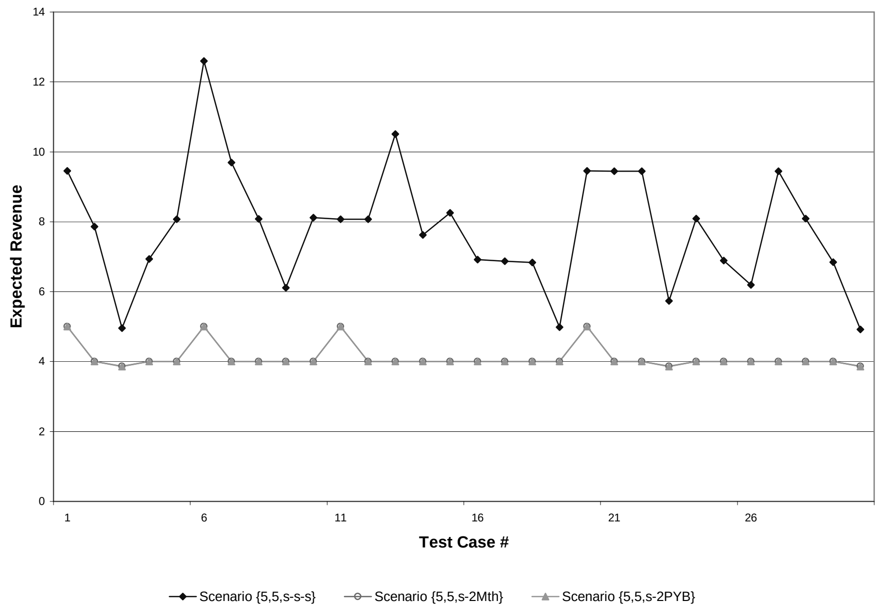

Figure 5.10: Comparison of the expected revenue among different auction scenarios when our agent plays its MCA strategy and the other agents' valuation are drawn from a uniform distribution.

sequential first-price, sealed-bid auctions is to bid the expected price of the object in each auction. This result is developed under the assumption that only the clearing price is revealed in previous auctions. In many current online auction environments, the actual bids and their associated bidders are revealed. In addition, we are not aware of any research on sequences of auctions with different rules.

Monte Carlo sampling has been previously used in conjunction with games of incomplete information. Frank et al. [21] describes an empirical study of the use of the Monte Carlo sampling method on a simple complete binary game tree. They draw the discouraging conclusion that the error rate quickly approaches 100% as the depth of the game increases. However, perhaps because Frank et al. consider only pure strategy equilibrium in a two-person, zero-sum game, these negative results did not evidence themselves in our study.

Howard James Bampton [3] investigated the use of Monte Carlo sampling to create a heuristic policy for the (imperfect information) game of Bridge. In Bampton's paper, he simply collected the player's decision in every sampled game and accumulated the chance-minimax values for each alternative at each decision node. Our method of accumulating sampled data is quite different from Bampton's approach, again because our game is not a two-player zero-sum game.

Researchers in artificial intelligence have recently been studying trading agents. A significant amount of work has gone into agents for the Trading Agent Competition (TAC) [29, 98, 107]. The TAC environment is significantly more complex than the simple scenarios presented here, and to date, none of the implemented agents model opponent behavior in a significant way.

Anthony, et al. [1] investigate agents that can participate in multiple online auctions. The authors posit a set of "tactics" and then empirically compare the performance of these tactics in a simulated market that consists of simultaneous and sequential English, Dutch, and Vickrey auctions. While the bidding strategies seem to resonate with particular aspects of human behavior (e.g., the "desperateness" strategy), they do not seem to have a foundation in any theory.

Boutilier et al. [7] develop a sequential auction model in which the agent values combinations of resources while all other participants value only a single item. Unlike our model, the Boutilier formulation does not explicitly model the opponents, though like our model it benefits from a dynamic programming approach to solving the decision problem.

Hon-Snir et al., [35] propose an iterative learning approach to solve repeated first-price auctions. They develop a repeated auction model which converges to an equilibrium strategy for a one-shot auction after many rounds of repeated auctions. In addition to the differences in overall structure of the marketplace, their work differs from ours in that they treat the other bidders as naive players. Specifically, they assume the opponents' next bid vectors are distributed according a weighted empirical distribution of their past bid vectors.

5.7 Conclusions

This study represents a first step in exploring the implementation of computational game theory in a simple trading agent. We show how Monte Carlo sampling can be used to construct a bidding policy that performs comparably to the subgame perfect equilibrium. This strategy takes advantage of information revealed in prior auctions in the sequence to improve play in later auctions. Importantly, the architecture is flexible, in that it can handle a variety of simple auction types, and different types of other bidders. Equally important, the approach is computationally limited by our ability to solve the component games, which suggests that algorithms for solving component games, particularly ones with well-structured payoff and action spaces, is an important area for further research.

We plan to continue this work and integrate more auction types, and to explore scenarios in which the agent's and other bidders' preferences are more complex, including scenarios in which the buyers may want more than one item. We would also like to add an aggregate buyer to the model to represent the large number of unmodeled opponents often found in public markets. Finally, we plan to explore auction sequences in which the bidders' valuations are correlated across the items, but not necessarily identical.

Chapter 6

The Non-Existence of Equilibrium in Sequential Auctions when Bids are Revealed

6.1 Introduction

A participant in a sequential auction must construct a strategy that is optimal for the sequence as a whole, and not just for an individual auction [19]. A natural approach is to model the sequential auctions as an extensive form game and solve for the equilibria. The outcome of this type of analysis depends upon critical assumptions in the model being studied.

Weber surveyed the research on sequential auctions and concluded that, with symmetric, risk-neutral bidders and identical items, the equilibrium price in a single-unit demand, first-price, sealed-bid sequential auction is a martingale. Weber's model examines two different price announcement schemes that enable the remaining bidders to infer the winning bidders' true valuation. The critical difference between Weber's model and ours is that we look at the case in which the auctioneer reveals all of the bids, not just the winner, at the end of the auction. Revealing all bids in an auction is popular on current public marketplaces such as eBay. As far as we know, none of the theoretical results have addressed the model with complete bid revelation.

Another model closely related to ours is that studied by Ortega-Reichert [86], in which two bidders bid on two items sold in a sequence of first-price, sealed-bid auctions. Ortega-Reichert derived equilibrium results for his model, and showed the signaling effects of the first bid on the second auction. However, his model differs from ours in a significant way that impacts the ability to derive a pure-strategy equilibrium. In the Ortega-Reichert model, the bidders have valuations for the two objects that are derived from a common distribution with an unknown parameter. The information revealed in the first auction affects each bidder's estimate of the value of the unknown parameter, and therefore their belief about their ability to win the second good. In our model, we consider a sequence of identical goods for which the bidders have a constant valuation. We show that a strategy that would reveal the bidders' valuations after the first auction would turn the remaining auctions into games of complete information.

Hausch [32] derived the necessary conditions for a symmetric equilibrium in Milgrom and Weber's general symmetric model by applying the signal game idea from Ortega-Reichert's model. Krishna [51] noted that in Weber's model the price quotes of the first period have no effect on the equilibrium bids in the second period. McAfee and Vincent [63] found a declining price pattern in symmetric sequential auctions when bidders have non-decreasing risk aversion. In another paper, the same authors [64] examined the equilibrium when a seller can post a reserve price in sequential auctions. Elmaghraby [15] studied the sequential second-price auction of heterogeneous items and concludes that the ordering of items effects the efficiency of the auction. Many other papers have addressed other variations of sequential auction models (e.g., [27, 36, 39]).

The remainder of the chapter proceeds as follows. In Section 6.2, we present a model of sequential auctions and point out the difference between Weber's model and our model. In Section 6.3 we discuss the symmetric equilibrium in Weber's model and show that Weber's equilibrium is not a solution to our model. In Section 6.3.3, we prove the non-existence of a symmetric pure-strategy equilibrium in the model for both first-price and second-price auctions. We also discuss the non-existence of asymmetric equilibrium in Section 6.4. We offer some conclusions in Section 6.5

6.2 The Model

There are K identical items for sale in a sequence of first-price, sealed-bid auctions. Exactly one item is sold in each auction. For convenience, we also use K to represent the set of items. There are N risk-neutral bidders, $N > K$, competing for the K items. Let A be the set of bidders. Each bidder has single-unit demand and will withdraw from the game once she wins one item. The bidders' valuations are independent observations of a nonnegative random variable, V, with a commonly known continuous cumulative distribution function, F, and its associated probability density function, f. We assume that F is continuous and differentiable in the domain of the valuation variables. Each bidder knows the value of the object to herself (the private values assumption), but not that of the other bidders.

Without loss of generality, we designate bidder 0 as the bidder whose strategy we are analyzing. Let $n = N - 1$, and let the other bidders be indexed from 1 to n. Let x be the true value of bidder 0's valuation and let y_i be the true value of bidder i's valuation, $i \in \{1, ..., n\}$. Without loss of generality, let Y_j be the $(n-j+1)$-st order statistic of $\{Y_1, ..., Y_n\}$. Thus, we have $Y_1 \leq Y_2 \leq ... \leq Y_n$. Let F_{Y_j} be the cdf of variable Y_j and let f_{Y_j} be the pdf of Y_j. F^k is the multiplication of F k times. Because bidders have identical information about each other's valuations at the beginning of the sequential auction, we refer to the model as the *symmetric sequential auction model*.

The key difference between our model and Weber's model [106] is the information revealed by the auctioneer. There are two different price quotes in Weber's model: the first announces only that an object has been sold, while the second announces also the sale price, p. Weber concludes that both price quotes yield the same equilibrium solution. We demonstrate that Weber's results do not hold if the auctioneer reveals *all bids* after the auction terminates.

Let $\beta_{k,i}(x)$ denote bidder i's bid function, which, given her valuation, x, the bidder can use to compute her bid, b_i, in auction k. As is common, we assume that these strategy functions are continuous, monotone and strictly increasing in the valuation. We also assume $\beta_k(x)$ is invertible, which means that a bidder's valuation can be inferred with certainty from the bid she makes [51, 106]. We assume $\beta(0) = 0$.

A *symmetric equilibrium* is an outcome in which all players adopt the same strategy. In our sequential auction model, a joint outcome is symmetric if $\beta_{k,i} = \beta_{k,j}$ for all bidders i and j. It is a symmetric equilibrium if no bidder can unilaterally increase her payoff by deviating from the symmetric strategy.

It has been shown that equilibrium does not exist in first price auction with continuous strategy space with complete information due to the discontinuity of the payoff function [53]. We address this technical issue using the technique proposed by Maskin and Riley [60]: a second round Vickrey auction is used to break the tie, if any. With the introduction of a second round Vickrey auction tie breaking rule, there exists a pure strategy equilibrium in which the highest type bidder bids a price equal to the second highest type bidder's valuation, and the other bidders bid at their true valuation. The introduction of this tie breaking is primarily a theoretical technicality because the probability of ties is zero when the strategy space is continuous.

6.3 Symmetric Equilibria in Sequential Auctions

6.3.1 Weber's Equilibrium

Weber derives a unique symmetric equilibrium for his model in which each bidder bids the expected value of the $(K+1)$th highest bidder assuming her own bid was the kth highest bid.[1] That is,

$$\beta_k(x) = E[Y_{N-K} | Y_{N-k} < x < Y_{N-k+1}]. \tag{6.1}$$

It is natural to question whether the price announcement in the first auction will influence the bidders' behaviors in the second auction. However, since the winner leaves the game, the remaining bidders have the same information about the rest of the game. A proof in [51] shows that the later period strategy is independent of the previous price announcement. As a result, for each bidder, the beliefs about the other bidder's valuation distributions remain unchanged. Weber explains that the type independence and symmetry assumptions make the equilibrium strategies independent of the two different price quotes [106].

However, when the auctioneer reveals all of the bids after each auction in the sequence, the above strategy is no longer an equilibrium strategy. The next section presents an example to demonstrate an individual bidder's incentive to deviate, and the following section proves the general case.

6.3.2 A Counter Example when Bids are Revealed

Consider a sequence of two auctions with bidders that follow the symmetric strategy in equation (6.1). That is, bidder i bids $b_i = \beta_{1,i}(v_i)$. Because β is invertible, after seeing the bids, every bidder can compute $v_i = \beta_1^{-1}(b_i)$, for all bidders remaining in the auction. As a result, the second auction becomes a game of complete information.

It is straightforward to show that a bidder can be better off by unilaterally deviating from equation (6.1) in the first auction.

[1] Interested readers may refer to [51], Section 15, for the proof.

Example 6.3.1. *Suppose there are $N = 10$ bidders in a sequence of two, first-price, sealed-bid auctions. If bidder 0 uses the strategy suggested by Weber, she will bid $\beta_1(x)$ in the first auction. Similarly, bidder i will bid $\beta_1(y_i)$. After the first auction, the second auction becomes a complete information auction. The second highest bidder wins the second item at the price of the third highest valuation. We have the following cases:*

1. *If $x > Y_9$, bidder 0 wins the first item and pays $E[Y_8|x > Y_9]$.*

2. *If $Y_9 > x > Y_8$, bidder 0 loses the first item, but will win the second auction with an expected payment of $E[Y_8|Y_9 > x > Y_8]$. The payoff is an expectation because bidder 0 will not obtain the true information about y_8 until the first auction completes.*

3. *If $Y_8 > x$, bidder 0 loses both the first and the second auctions.*

Now, suppose bidder 0 deviates from $\beta_1(x)$ to $\beta_1^{'}(x) = 0$, while the other bidders stick to Weber's strategy. After the first auction, the other bidders infer that bidder 0's valuation is 0, and is therefore not a factor in their decisions. Although Bidder 0 will always lose the first auction, she can benefit from this deception, as evidenced by the following four exhaustive cases.

1. *If $x > Y_9$, Y_9 wins the first item. In the second auction, all of the other bidders will believe that Y_8 has the highest valuation and Y_7 is the second highest. Thus, bidder 8 will bid y_7. On average, bidder 0 will be able to win the second item at price $E[Y_7|x > Y_9]$. Since $E[Y_8|x > Y_9] \geq E[Y_7|x > Y_9]$, $x - E[Y_7|x > Y_9] \geq x - E[Y_8|x > Y_9]$. Thus, bidder 0 will be better off by deviating in this case.*

2. *If $Y_9 > x > Y_8$, bidder 0 will again win the second auction at $E[Y_7|Y_9 > x > Y_8]$.*

	Utility When Using Weber's Strategy	Utility When Deviating to Strategy $z = 0$	Increase in Utility When Deviating
$x > Y_9$	$x - E[Y_8 \mid x > Y_9]$	$x - E[Y_7 \mid x > Y_9]$	> 0
$Y_9 > x > Y_8$	$x - E[Y_8 \mid Y_9 > x > Y_8]$	$x - E[Y_7 \mid Y_9 > x > Y_8]$	> 0
$Y_8 > x > Y_7$	0	$x - E[Y_7 \mid Y_8 > x > Y_7] - \epsilon$	> 0
$Y_7 > x$	0	0	$= 0$

Table 6.1: The expected utility of bidder 0 in the sequential FPSB auctions.

3. *If $Y_8 > x > Y_7$, bidder 8 will believe that the second highest valuation in the second auction is Y_7, and will bid at y_7. Again, bidder 0 can bid at $y_7 + \epsilon$ and win the second item, where the ϵ term is included to avoid the tie with Y_8. Thus, bidder 0 will expect to pay $E[Y_7|Y_8 > x > Y_7] - \epsilon$.*

4. *If $Y_7 > x$, bidder 0 will lose both items.*

Thus, bidder 0 will have a greater expected payoff by unilaterally deviating in the first auction. A comparison of the cases is shown in Table 6.1.

Example 6.3.1 demonstrates that a bidder would be better off by unilaterally deviating from Weber's strategy when the other bidders use Weber's strategies in the first auction. Thus, Weber's equilibrium strategies cannot be an equilibrium in this new model of sequential auctions. We now address the question of whether any pure strategy, symmetric equilibrium exists in this game.

6.3.3 Non-Existence of Symmetric Equilibrium in Sequential First-Price Auctions

We now address the general question of the existence of symmetric equilibrium in the sequential auction with bid revelation. We continue to consider a sequence of two auctions. Assume that there exists a symmetric pure strategy equilibrium, such that every bidder uses the same strategic function, β, in the first auction. With this assumption, and the previous assumption that β is strictly monotonically increasing and invertible, every bidder can infer every other bidder's true valuation after the first auction. As a result, the second and future auctions become complete information games.

Thus, we restrict our analysis to the first auction, in which the bidders have incomplete information. The definition of a symmetric equilibrium requires that a bidder cannot be better off by unilaterally deviating from β when all other bidders are playing β. Let $u(x)$ denote the payoff to a bidder if she bids $\beta(x)$. Let $u(x, z|z \geq x)$ denote the payoff of this bidder if she deviates from $\beta(x)$ to a higher bid $\beta(z)$. Similarly, we let $u(x, z|z \leq x)$ denotes the payoff of this bidder if she deviates from $\beta(x)$ to a lower bid $\beta(z)$.

When $z \geq x$, bidder 0 will win the first item if z is larger than the highest valuation of the other bidders. Otherwise, she will win the second item if x is larger than the second highest valuation of the other bidders. Bidder 0's payoff function can be written as [2]

$$\begin{aligned} u(x, z|z \geq x) \quad &= \quad \Pr(Y_n < z)[x - \beta(z)] \\ &\quad + \Pr(Y_{n-1} < x \leq z < Y_n) \times \\ &\quad\; [x - E[Y_{n-1}|Y_{n-1} < x \leq z < Y_n]]. \end{aligned} \tag{6.2}$$

The first term results from the event $Y_n < z$ in which bidder 0 wins the first auction. The second term captures the case that bidder 0 loses the first auction and wins the second by bidding the revealed value of the third highest bidder.

We now consider the case where $z \leq x$. If z is larger than the highest valuation of the other bidders, bidder 0 will win the first item. Otherwise, bidder 0 may still be able to win the second item, depending upon the revealed valuations of the other bidders. In the following analysis, bidder 0 may bid against a bidder who has a type greater than bidder 0's. In such a case, bidder 0 would lose the tie-breaker unless she bids slightly above the expected bid of the bidder with the higher type. We introduce the small value, ϵ, which bidder 0 uses to avoid a tie with a higher type bidder.

There are four variations of non-zero outcomes:

- Case 1: $Y_{n-1} < z < Y_n$. Bidder 0 loses the first item; however, both z and x are larger than the second highest valuation of the other bidders such that bidder 0 will win the second item and expect to pay $E[Y_{n-1}|Y_{n-1} < z < Y_n]$.

[2] When the strategy space and the probability function are continuous, the probability that two bidders tie at the same valuation or strategy is zero. Because of this, we will ignore edge equalities in the following equations.

- Case 2: $Y_{n-2} < z < Y_{n-1}$. In this case, bidder 0 still loses the first item. However, the bidder with the second highest valuation infers from the first auction that the third highest valuation is z. As a result, he will bid z in the second auction. Bidder 0 can bid $z + \epsilon$ to outbid the bidder with rank Y_{n-1}.

- Case 3: $z < Y_{n-2} < x < Y_{n-1}$. In this case, the bidder with the second highest valuation will bid what appears to be the third highest value: $E[Y_{n-2}|z < Y_{n-2} < x < Y_{n-1}]$. Bidder 0 needs to bid $E[Y_{n-2}|z < Y_{n-2} < x < Y_{n-1}] + \epsilon$ to outbid the bidder with the second highest type.

- Case 4: $z < Y_{n-2} < Y_{n-1} < x$. This case is the same as the above case with the exception that bidder 0 will win the tie breaker and so does not need to add ϵ to her bid to win the second item.

Thus, when $z \leq x$, the payoff function can be written as

$$\begin{aligned} u(x, z, \epsilon|z \leq x) \quad &= \quad \Pr(Y_n < z)[x - \beta(z)] \\ &\quad + \Pr(Y_{n-1} < z < Y_n)[x - E[Y_{n-1}|Y_{n-1} < z < Y_n]] \\ &\quad + \Pr(Y_{n-2} < z < Y_{n-1})[x - (z + \epsilon)] \\ &\quad + \Pr(z < Y_{n-2} < x < Y_{n-1}) \times \\ &\qquad [x - (E[Y_{n-2}|z < Y_{n-2} < x < Y_{n-1}] + \epsilon)] \\ &\quad + \Pr(z < Y_{n-2} < Y_{n-1} < x) \times \\ &\qquad [x - E[Y_{n-2}|z < Y_{n-2} < Y_{n-1} < x]]. \end{aligned}$$

In the above equation, the first term represents the case when bidder 0 wins the first auction. The next four terms represent the cases 1–4 above.

As ϵ goes to zero, $u(x, z, \epsilon|z \leq x)$ asymptotically goes to

$$\begin{aligned} u(x, z|z \leq x) \quad &= \quad \Pr(Y_n < z)[x - \beta(z)] \\ &\quad + \Pr(Y_{n-1} < z < Y_n)[x - E[Y_{n-1}|Y_{n-1} < z < Y_n]] \\ &\quad + \Pr(Y_{n-2} < z < Y_{n-1})[x - z] \\ &\quad + \Pr(z < Y_{n-2} < x)[x - E[Y_{n-2}|z < Y_{n-2} < x]]. \end{aligned} \tag{6.3}$$

For $\beta(x)$ to be bidder 0's best response to the other bidders playing β, it must be true that $u(x,z|z \geq x) \leq u(x)$, and $u(x,z,\epsilon|z \leq x) \leq u(x)$. As ϵ goes to zero, the symmetric equilibrium requires

$$\begin{aligned} & u(x,z|z \geq x) \leq u(x), \text{and} && (6.4) \\ & u(x,z|z \leq x) \leq u(x). \end{aligned}$$

It follows from equations (6.2) and (6.3) that $u(x,z|z \leq x)$ and $u(x,z|z \geq x)$ are continuous and differentiable because the probability functions are continuous and differentiable. Also, we know $u(x,z|z \leq x) = u(x,z|z \geq x) = u(x)$ when $z = x$. We now present our main result.

Theorem 6.3.1. *In the symmetric sequential auction model with full bid revelation, there does not exist a symmetric pure-strategy equilibrium.*

Proof: We prove the result by contradiction. We first assume that in the symmetric sequential auction model in which all bids are revealed, there exists a symmetric pure-strategy equilibrium, β.

When $z \geq x$, we refer to it as a right hand side (RHS) deviation. Similarly, $z \leq x$ is a left hand side (LHS) deviation. In the following discussion, we replace β with β_{RHS} in $u(x,z|z \geq x)$ and replace β with β_{LHS} in $u(x,z|z \leq x)$. Our target is to solve β_{RHS} and β_{LHS} respectively from equations (6.2) and (6.3). From the assumption that β is a symmetric pure strategy that is continuous, monotonically increasing, and invertible, it follows that $\beta_{\text{RHS}}(x) = \beta_{\text{LHS}}(x)$.

We can rewrite equation (6.4) as

$$\begin{aligned} & \frac{u(x,z|z \geq x) - u(x)}{z - x} \leq 0, \text{for all } z \geq x, \text{ and} \\ & \frac{u(x,z|z \leq x) - u(x)}{z - x} \geq 0, \text{for all } z \leq x. \end{aligned}$$

As a result, there must exist a small enough deviation $|z - x|$ such that taking the first order condition on both sides of x gives

$$\frac{\partial u(x,z|z \geq x)}{\partial z} \leq 0, \text{for all } z \geq x, \text{and} \tag{6.5}$$

$$\frac{\partial u(x,z|z \leq x)}{\partial z} \geq 0, \text{for all } z \leq x. \tag{6.6}$$

Equation (6.5) can be derived from equation (6.2). When $z \geq x$, we have

$$
\begin{aligned}
&u(x,z|z\geq x)\\
&\quad = F_{Y_n}(z)[x-\beta_{\text{RHS}}(z)]\\
&\qquad + n(1-F(z))F^{n-1}(x)\times\\
&\qquad\quad [x-\frac{\int_{-\infty}^{x}\int_{z}^{+\infty} y_{n-1}f_{y_{n-1},y_n}(y_{n-1},y_n)dy_n dy_{n-1}}{n(1-F(z))F^{n-1}(x)}]\\
&\quad = F_{Y_n}(z)[x-\beta_{\text{RHS}}(z)] + n(1-F(z))F^{n-1}(x)x\\
&\qquad - \int_{-\infty}^{x}\int_{z}^{+\infty} y_{n-1}n(n-1)f(y_{n-1})f(y_n)F^{n-2}(y_{n-1})dy_n dy_{n-1}\\
&\quad = F_{Y_n}(z)[x-\beta_{\text{RHS}}(z)] + n(1-F(z))F^{n-1}(x)x\\
&\qquad - \int_{-\infty}^{x} y_{n-1}n(n-1)f(y_{n-1})F^{n-2}(y_{n-1})dy_{n-1}[1-F(z)].
\end{aligned}
$$

Solving the first order condition, we obtain

$$
\begin{aligned}
\frac{\partial u(x,z|z\geq x)}{\partial z} &= f_{Y_n}(z)[x-\beta_{\text{RHS}}(z)] - F_{Y_n}(z)\beta'_{\text{RHS}}(z) - nf(z)F^{n-1}(x)x\\
&\quad + \int_{-\infty}^{x} y_{n-1}n(n-1)f(y_{n-1})F^{n-2}(y_{n-1})dy_{n-1}f(z)\\
&\leq 0.
\end{aligned}
$$

From the definition of equilibrium, we know that $u(x,z|z\geq x)$ is maximized at $z=x$. Because $F_n(x)=F^n(x)$ and $f_{Y_n}(x)=nf(x)F^{n-1}(x)$, setting $z=x$ allows us to reduce the above equation to

$$
\begin{aligned}
\frac{\partial u(x,z|z\geq x)}{\partial z}|_{z=x} &= -[F_{Y_n}(x)\beta_{\text{RHS}}(x)]'\\
&\quad + \int_{-\infty}^{x} y_{n-1}n(n-1)f(y_{n-1})F^{n-2}(y_{n-1})dy_{n-1}f(x)\\
&= 0.
\end{aligned}
$$

As a result,

$$
[F_{Y_n}(x)\beta_{\text{RHS}}(x)]' = \int_{-\infty}^{x} n(n-1)y_{n-1}f(y_{n-1})F^{n-2}(y_{n-1})dy_{n-1}f(x). \tag{6.7}
$$

Now, let us consider the case where $z \leq x$. Equation (6.3) can be rewritten as

$$
\begin{aligned}
&u(x,z|z\leq x)\\
&= F_{Y_n}(z)[x-\beta_{\text{LHS}}(z)] + n(1-F(z))F^{n-1}(z)x\\
&\quad -\int_{-\infty}^{z}\int_{z}^{+\infty} y_{n-1}n(n-1)f(y_n)f(y_{n-1})F^{n-2}(y_{n-1})dy_n dy_{n-1}\\
&\quad +\int_{-\infty}^{z}\int_{z}^{+\infty} n(n-1)(n-2)F(y_{n-2})^{n-3}f(y_{n-2})f(y_{n-1})\times\\
&\qquad [1-F(y_{n-1})]dy_{n-1}dy_{n-2}[x-z]\\
&\quad +x\int_{z}^{x} f_{Y_{n-2}}(y_{n-2})dy_{n-2} - \int_{z}^{x} y_{n-2}f_{Y_{n-2}}(y_{n-2})dy_{n-2}\\
&= F_{Y_n}(z)[x-\beta_{\text{LHS}}(z)] + n(1-F(z))F^{n-1}(z)x\\
&\quad -\int_{-\infty}^{z} n(n-1)y_{n-1}f(y_{n-1})F^{n-2}(y_{n-1})dy_{n-1}[1-F(z)]\\
&\quad +n(n-1)(n-2)\int_{-\infty}^{z} F(y_{n-2})^{n-3}f(y_{n-2})dy_{n-2}\times\\
&\qquad \int_{z}^{+\infty} f(y_{n-1})[1-F(y_{n-1})]dy_{n-1}[x-z]\\
&\quad +x\int_{z}^{x} f_{Y_{n-2}}(y_{n-2})dy_{n-2} - \int_{z}^{x} y_{n-2}f_{Y_{n-2}}(y_{n-2})dy_{n-2}\\
&= F_{Y_n}(z)[x-\beta_{\text{LHS}}(z)] + n(1-F(z))F^{n-1}(z)x\\
&\quad -\int_{-\infty}^{z} n(n-1)y_{n-1}f(y_{n-1})F^{n-2}(y_{n-1})dy_{n-1}[1-F(z)]\\
&\quad +\frac{n(n-1)}{2}F^{n-2}(z)[1-F(z)]^2[x-z]\\
&\quad +x\int_{z}^{x} f_{Y_{n-2}}(y_{n-2})dy_{n-2} - \int_{z}^{x} y_{n-2}f_{Y_{n-2}}(y_{n-2})dy_{n-2}.
\end{aligned}
$$

The first order condition is

$$\begin{aligned}\frac{\partial u(x,z|z\le x)}{\partial z} &= f_{Y_n}(z)[x-\beta_{\text{LHS}}(z)] - F_{Y_n}(z)\beta'_{LHS}(z)\\ &\quad +n(n-1)F^{n-2}(z)f(z)[1-F(z)]x - nF^{n-1}(z)f(z)x\\ &\quad -n(n-1)zf(z)F^{n-2}(z)[1-F(z)]\\ &\quad +\int_{-\infty}^{z} n(n-1)y_{n-1}f(y_{n-1})F^{n-2}(y_{n-1})dy_{n-1}f(z)\\ &\quad +\frac{n(n-1)(n-2)}{2}F^{n-3}(z)f(z)[1-F(z)]^2[x-z]\\ &\quad -n(n-1)F^{n-2}(z)[1-F(z)]f(z)[x-z]\\ &\quad -\frac{n(n-1)}{2}F^{n-2}(z)[1-F(z)]^2\\ &\quad -f_{Y_{n-2}}(z)x + zf_{Y_{n-2}}(z)\\ &\ge 0.\end{aligned}$$

At $z=x$,

$$\begin{aligned}&\frac{\partial u(x,z|z\le x)}{\partial z}|_{z=x}\\ &= f_{Y_n}(x)[x-\beta_{\text{LHS}}(x)] - F_{Y_n}(x)\beta'_{\text{LHS}}(x) - nF^{n-1}(x)f(x)x\\ &\quad +\int_{-\infty}^{x} n(n-1)y_{n-1}f(y_{n-1})F^{n-2}(y_{n-1})dy_{n-1}f(x)\\ &\quad -\frac{n(n-1)}{2}F^{n-2}(x)[1-F(x)]^2\\ &= -[F_{Y_n}(x)\beta_{\text{LHS}}(x)]'\\ &\quad +\int_{-\infty}^{x} n(n-1)y_{n-1}f(y_{n-1})F^{n-2}(y_{n-1})dy_{n-1}f(x)\\ &\quad -\frac{n(n-1)}{2}F^{n-2}(x)[1-F(x)]^2\\ &\ge 0.\end{aligned}$$

Thus, we have

$$\begin{aligned}[F_{Y_n}(x)\beta_{\text{LHS}}(x)]' &\le \int_{-\infty}^{x} n(n-1)y_{n-1}f(y_{n-1})F^{n-2}(y_{n-1})dy_{n-1}f(x)\\ &\quad -\frac{n(n-1)}{2}F^{n-2}(x)[1-F(x)]^2. \qquad (6.8)\end{aligned}$$

We have now obtained a closed form solution for computing both $\beta_{\text{RHS}}(x)$ and $\beta_{\text{LHS}}(x)$.

Combining equations (6.7) and (6.8), and noting that $\frac{n(n-1)}{2}F^{n-2}(x)[1-F(x)]^2 > 0$, we see that

$$\begin{aligned}
[F_{Y_n}(x)\beta_{\text{LHS}}(x)]' &= \int_{-\infty}^{x} n(n-1)y_{n-1}f(y_{n-1})F^{n-2}(y_{n-1})dy_{n-1}f(x) \\
&\quad -\frac{n(n-1)}{2}F^{n-2}(x)[1-F(x)]^2 \\
&< \int_{-\infty}^{x} n(n-1)y_{n-1}f(y_{n-1})F^{n-2}(y_{n-1})dy_{n-1}f(x) \\
&= [F_{Y_n}(x)\beta_{\text{RHS}}(x)]'.
\end{aligned}$$

Because $\beta_{\text{RHS}}(0) = \beta_{\text{LHS}}(0) = \beta(0) = 0$, $F_{Y_n}(0)\beta_{\text{RHS}}(0) = F_{Y_n}(0)\beta_{\text{LHS}}(0) = 0$. By integrating both sides of the above equation, we obtain for all x

$$\beta_{LHS}(x) < \beta_{RHS}(x), \tag{6.9}$$

which implies that β does not exist for any x. This result contradicts the assumption that β is monotonically increasing and continuous, having $\beta_{LHS}(x) = \beta_{RHS}(x)$. Thus, in the symmetric, sequential auction model with all bids revealed, there does not exist a symmetric pure-strategy equilibrium. $\diamond$

It is worth noting that $\beta_{LHS}(x) < \beta_{RHS}(x)$ does not imply that there exist two different equilibrium strategy functions. The whole proof process shows that there does not exist a definition for β at any specified point x because the asymptotic limits from either side are not equal.

6.3.4 Non-Existence of Symmetric Equilibrium in Sequential Vickrey Auctions

Weber characterized the equilibrium in the sequential, Vickrey auction scenario with only the winner's bid revealed as for all $k < K$

$$\beta_k(x) = E[Y_{N-K}|Y_{N-k-1} < x < Y_{N-k}]. \tag{6.10}$$

In the last auction, where $k = K$, each bidder bids her true valuation. In the auctions prior to the last, each bidder bids her expectation of the ($K+1$)-th highest bidder assuming that she is at or above K.

To analyze the sequential Vickrey model with all bids revealed, we again assume that there exists a symmetric, pure-strategy equilibrium such that every bidder uses the same strictly monotone increasing and invertible bidding function, β, in the first auction. Thus, every bidder can infer every other bidder's true valuation after the first auction, and the second and future auctions become complete information games.

	Utility when Using Weber's Strategy	Utility Deviating to $z = 0$	Increase in Utility when Deviating
$x > Y_3$	$x - E[Y_2 \vert Y_3 > x > Y_2]$	x	> 0
$Y_3 > x > Y_2$	$x - E[Y_2 \vert Y_2 > x > Y_1]$	x	> 0
$Y_2 > x > Y_1$	$x - E[Y_1 \vert Y_2 > x > Y_1]$	x	> 0
$Y_1 > x$	0	x	> 0

Table 6.2: The expected utility of bidder 0 in the sequential Vickrey auctions.

In a sequence of $K-1$ Vickrey auctions with complete information, it is a Nash equilibrium for the top $K-1$ players to bid at the K-th highest remaining valuation, while all others bid their true valuation. This conclusion, however, leads to the observation that Weber's strategy is not an equilibrium in the first auction when the bids of all bidders are revealed. As in Section 6.3.3, a 3-item, 4-bidder sequence of Vickrey auctions illustrates how one bidder can improve her expected utility by misrepresenting her valuation in the first auction. The four conditions and their expected payoffs are shown in Table 6.2.

We now examine a sequence of three[3] Vickrey auctions with an arbitrary number of bidders. We concentrate our analysis on the first auction, in which the bidders have incomplete information, and show that a symmetric, pure-strategy equilibrium does not exist.

When all bidders play β and Bidder 0 selects $z \geq x$, she will win the first item if z is larger than the highest valuation of the other bidders, and will pay the second highest bid, β_{Y_n}. Otherwise, she will win the second item if x is larger than the second highest valuation of the other bidders. Bidder 0's payoff function can be written as

$$\begin{aligned} &u(x, z|z \geq x) \qquad (6.11)\\ &\quad = \Pr(Y_n < z)[x - E[\beta(Y_n)|Y_n < z]] \\ &\quad\quad + \Pr(Y_{n-2} < x \leq z < Y_n)[x - E[Y_{n-2}|Y_{n-2} < x \leq z < Y_n]]. \end{aligned}$$

The first term results from the event $Y_n < z$ in which bidder 0 wins the first auction. The second term captures the case in which bidder 0 loses the first auction and wins the second by bidding the revealed value of the fourth highest bidder.

We now consider the case where bidder 0 chooses $z \leq x$. If z is larger than the highest valuation of the other bidders, bidder 0 will win the first item. Otherwise, bidder 0 may still be able

[3]The extension to an arbitrary number of auctions follows easily from the three auction case.

to win the second or third item, depending upon the revealed valuations of the other bidders.

There are four non-zero outcomes:

- Case 1: $Y_{n-2} < z < Y_n$. Bidder 0 loses the first item; however, both z and x are larger than the third highest valuation among the other bidders and bidder 0 will be able to win one of the next two auctions and expect to pay $E[Y_{n-2}|Y_{n-2} < z < Y_n]$.
- Case 2: $Y_{n-3} < z < Y_{n-2}$. In this case, bidder 0 still loses the first item. However, the bidders with the second and third highest valuations infer from the first auction that the fourth highest valuation is z and, as a result, will bid z in the second auction. This allows bidder 0 to outbid bidder Y_{n-2} by bidding $z + \epsilon$, thus stealing an item when she is not one of the three highest valuing bidders.
- Case 3: $z < Y_{n-3} < x$. In this case, the bidder with the third highest valuation will bid what appears to be the fourth highest value. There are two sub-cases. The first sub-case is $x < Y_{n-2}$, so Bidder 0 needs to bid $E[Y_{n-2}|z < Y_{n-3} < x < Y_{n-2}] + \epsilon$ to outbid the bidder with the third highest type. The second sub-case is $z < Y_{n-3} < Y_{n-2} < x$, so Bidder 0 does not need to add ϵ to her bid to win the second/third item.

As ϵ goes to zero, asymptotically, the payoff function can be written as

$$\begin{aligned} u(x, z|z \leq x) &= \Pr(Y_n < z)[x - E[\beta(Y_n)|Y_n < z]] \\ &\quad + \Pr(Y_{n-2} < z < Y_n)[x - E[Y_{n-2}|Y_{n-2} < z < Y_n]] \\ &\quad + \Pr(Y_{n-3} < z < Y_{n-2})[x - z] \\ &\quad + \Pr(z < Y_{n-3} < x)[x - E[Y_{n-3}|z < Y_{n-3} < x]]. \end{aligned} \tag{6.12}$$

As a condition of symmetric equilibrium, equation 6.4 is also true for sequential Vickrey auctions. We now present our second main result.

Theorem 6.3.2. *In a sequence of symmetric Vickrey auctions with full bid revelation, there does not exist a symmetric, pure-strategy equilibrium.*

Proof: We prove the result by contradiction following the same strategy as in the FPSB model. We first assume that there exists a symmetric pure-strategy equilibrium, β.

Again, we refer to the case where the bidder selects $z \geq x$ as a right hand side (RHS) deviation. Similarly, $z \leq x$ is a left hand side (LHS) deviation. In the following discussion, we replace β with β_{RHS} when $u(x,z|z \geq x)$, and replace β with β_{LHS} when $u(x,z|z \leq x)$. Furthermore, we replace $E[\beta(Y_n)|Y_n < z]$ with Θ_{RHS} in $u(x,z|z \geq x)$ and replace $E[\beta(Y_n)|Y_n < z]$ with β_{LHS} in $u(x,z|z \leq x)$. Our target is to solve β_{RHS} and β_{LHS} respectively from equations (6.11) and (6.12). Because we assume that β is continuous, monotonically increasing, and invertible, we should find that $\beta_{\text{RHS}}(x) = \beta_{\text{LHS}}(x)$ and $\Theta_{\text{RHS}}(x) = \Theta_{\text{LHS}}(x)$.

We can rewrite equation (6.11) as follows.

$$
\begin{aligned}
&u(x,z|z \geq x) \\
&\quad= F_{Y_n}(z)[x - \Theta_{\text{RHS}}(z)] \\
&\qquad+ \int_{-\infty}^{x} \int_{z}^{+\infty} f_{y_{n-2},y_n}(y_{n-2},y_n) dy_n dy_{n-2} \times \\
&\qquad\ [x - \frac{\int_{-\infty}^{x} \int_{z}^{+\infty} y_{n-2} f_{y_{n-1},y_n}(y_{n-2},y_n) dy_n dy_{n-2}}{\int_{-\infty}^{x} \int_{z}^{+\infty} f_{y_{n-2},y_n}(y_{n-2},y_n) dy_n dy_{n-2}}] \\
&\quad= F_{Y_n}(z)[x - \Theta_{\text{RHS}}(z)] \\
&\qquad+ \int_{-\infty}^{x} \int_{z}^{+\infty} [n(n-1)(n-2) f(y_{n-2}) f(y_n) \times \\
&\qquad\ F^{n-3}(y_{n-2})[F(y_n) - F(y_{n-2})]] dy_n dy_{n-2} x \\
&\qquad- \int_{-\infty}^{x} \int_{z}^{+\infty} [n(n-1)(n-2) y_{n-2} f(y_{n-2}) f(y_n) \times \\
&\qquad\ F^{n-3}(y_{n-2})[F(y_n) - F(y_{n-2})]] dy_n dy_{n-2} \\
&\quad= F_{Y_n}(z)[x - \Theta_{\text{RHS}}(z)] \\
&\qquad+ n(n-1)(n-2) F^{n-2}(x)[1 - F(z)] [\frac{1 + F(z)}{2(n-2)} - \frac{F(x)}{n-1}] x \\
&\qquad- \int_{-\infty}^{x} n(n-1)(n-2) y_{n-2} f(y_{n-2}) F^{n-3}(y_{n-2}) dy_{n-2} \frac{1 - F^2(z)}{2} \\
&\qquad+ \int_{-\infty}^{x} n(n-1)(n-2) y_{n-2} f(y_{n-2}) F^{n-2}(y_{n-2}) dy_{n-2} [1 - F(z)].
\end{aligned}
$$

Solving the first order condition, we obtain

$$
\begin{aligned}
&\frac{\partial u(x,z|z\geq x)}{\partial z}\\
&\quad = f_{Y_n}(z)[x-\Theta_{\text{RHS}}(z)] - F_{Y_n}(z)\Theta^{'}_{\text{RHS}}(z)\\
&\qquad +n(n-1)(n-2)F^{n-2}(x)\big[\frac{-2F(z)f(z)}{2(n-2)}+\frac{F(x)f(z)}{n-1}\big]x\\
&\qquad -\int_{-\infty}^{x} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-3}(y_{n-2})dy_{n-2}\times\\
&\qquad\quad \frac{[-2F(z)f(z)]}{2}\\
&\qquad +\int_{-\infty}^{x} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-2}(y_{n-2})dy_{n-2}\times\\
&\qquad\quad [-f(z)].\\
&\quad = -[F_{Y_n}(z)\Theta_{\text{RHS}}(z)]^{'} + nf(z)F^{n-1}(z)x\\
&\qquad +n(n-1)(n-2)F^{n-2}(x)\big[\frac{-2F(z)f(z)}{2(n-2)}+\frac{F(x)f(z)}{n-1}\big]x\\
&\qquad +F(z)f(z)\int_{-\infty}^{x} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-3}(y_{n-2})dy_{n-2}\\
&\qquad -f(z)\int_{-\infty}^{x} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-2}(y_{n-2})dy_{n-2}\\
&\quad = -[F_{Y_n}(z)\Theta_{\text{RHS}}(z)]^{'} + nf(z)F^{n-1}(z)x\\
&\qquad -n(n-1)(n-2)F^{n-2}(x)\big[\frac{F(z)}{n-2}-\frac{F(x)}{n-1}\big]f(z)x\\
&\qquad +\int_{-\infty}^{x} [n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-3}(y_{n-2})\times\\
&\qquad\quad [F(z)-F(y_{n-2})]]dy_{n-2}f(z)\\
&\qquad \leq 0.
\end{aligned}
$$

From the definition of equilibrium, we know that $u(x,z|z\geq x)$ reaches its optimal point when $z=x$. By setting $z=x$, the above equation reduces to

$$
\begin{aligned}
&\frac{\partial u(x,z|z\geq x)}{\partial z}|_{z=x}\\
&\quad = -[F_{Y_n}(x)\Theta_{\text{RHS}}(x)]^{'}\\
&\qquad +\int_{-\infty}^{x} n(n-1)(n-2)y_{n-2}f(y_{n-2})\times\\
&\qquad\quad F^{n-3}(y_{n-2})[F(x)-F(y_{n-2})]dy_{n-2}f(x)\\
&\quad = 0.
\end{aligned}
$$

As a result,

$$
\begin{aligned}
F_{Y_n}(x)\Theta_{\text{RHS}}(x)]' \\
&= \int_{-\infty}^{x} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-3}(y_{n-2}) \times \\
&\quad [F(x) - F(y_{n-2})]dy_{n-2}f(x).
\end{aligned}
\tag{6.13}
$$

Now, let us consider the case where $z \leq x$. Equation (6.12) can be rewritten as

$$
\begin{aligned}
&u(x, z|z \leq x) \\
&\quad = F_{Y_n}(z)[x - \Theta_{\text{LHS}}(z)] \\
&\qquad + \int_{-\infty}^{z}\int_{z}^{+\infty} f_{y_{n-2},y_n}(y_{n-2}, y_n)dy_n dy_{n-2} \times \\
&\qquad [x - \frac{\int_{-\infty}^{z}\int_{z}^{+\infty} y_{n-2}f_{y_{n-1},y_n}(y_{n-2}, y_n)dy_n dy_{n-2}}{\int_{-\infty}^{x}\int_{z}^{+\infty} f_{y_{n-2},y_n}(y_{n-2}, y_n)dy_n dy_{n-2}}] \\
&\qquad + \int_{-\infty}^{z}\int_{z}^{+\infty} \frac{n!}{(n-4)!2!} f(y_{n-3})f(y_{n-2})F^{n-4}(y_{n-3}) \times \\
&\qquad [1 - F(y_{n-2})]^2 dy_{n-3}dy_{n-2}(x - z) \\
&\qquad + \int_{z}^{x} \frac{n!}{(n-4)!3!} f(y_{n-3})F^{n-4}(y_{n-3})[1 - F(y_{n-2})]^3 dy_{n-3}x \\
&\qquad - \int_{z}^{x} \frac{n!}{(n-4)!3!} y_{n-3}f(y_{n-3})F^{n-4}(y_{n-3})[1 - F(y_{n-2})]^3 dy_{n-3}.
\end{aligned}
$$

$$
\begin{aligned}
&u(x,z|z\le x)\\
&\quad= F_{Y_n}(z)[x-\Theta_{\text{LHS}}(z)]\\
&\qquad+\int_{-\infty}^{z}\int_{z}^{+\infty} n(n-1)(n-2)f(y_{n-2})f(y_n)F^{n-3}(y_{n-2})\times\\
&\qquad\ [F(y_n)-F(y_{n-2})]dy_n dy_{n-2}x\\
&\qquad-\int_{-\infty}^{z}\int_{z}^{+\infty} n(n-1)(n-2)y_{n-2}f(y_{n-2})f(y_n)\times\\
&\qquad\ F^{n-3}(y_{n-2})[F(y_n)-F(y_{n-2})]dy_n dy_{n-2}\\
&\qquad+\frac{n!}{(n-4)!2!}\int_{-\infty}^{z} f(y_{n-3})F^{n-4}(y_{n-3})dy_{n-3}\times\\
&\qquad\ \int_{z}^{+\infty} f(y_{n-2})[1-F(y_{n-2})]^2 dy_{n-2}(x-z)\\
&\qquad+\frac{n!}{(n-4)!3!}\int_{z}^{x}\big[F^{n-4}(y_{n-3})-3F^{n-3}(y_{n-3})\times\\
&\qquad\ +3F^{n-2}(y_{n-3})-F^{n-1}(y_{n-3})\big]dF(y_{n-3})x\\
&\qquad-\int_{z}^{x}\frac{n!}{(n-4)!3!}y_{n-3}f(y_{n-3})F^{n-4}(y_{n-3})[1-F(y_{n-2})]^3 dy_{n-3}\\
&\quad= F_{Y_n}(z)[x-\Theta_{\text{LHS}}(z)]\\
&\qquad+n(n-1)(n-2)F^{n-2}(z)[1-F(z)]\big[\frac{1+F(z)}{2(n-2)}-\frac{F(z)}{n-1}\big]x\\
&\qquad-\int_{-\infty}^{z} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-3}(y_{n-2})dy_{n-2}\frac{1-F^2(z)}{2}\\
&\qquad+\int_{-\infty}^{z} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-2}(y_{n-2})dy_{n-2}[1-F(z)]\\
&\qquad+\frac{n!}{(n-3)!3!}F^{n-3}(z)[1-F(z)]^3(x-z)\\
&\qquad+\frac{n!}{(n-4)!3!}\big[\frac{F^{n-3}(x)-F^{n-3}(z)}{n-3}-\frac{3[F^{n-2}(x)-F^{n-2}(z)]}{n-2}+\\
&\qquad\ \frac{3[F^{n-1}(x)-F^{n-1}(z)]}{n-1}-\frac{F^{n}(x)-F^{n}(z)}{n}\big]x\\
&\qquad-\int_{z}^{x}\frac{n!}{(n-4)!3!}y_{n-3}f(y_{n-3})F^{n-4}(y_{n-3})[1-F(y_{n-2})]^3 dy_{n-3}.
\end{aligned}
$$

The first order condition is

$$
\begin{aligned}
&\frac{\partial u(x,z|z\le x)}{\partial z}\\
&\quad = -[F_{Y_n}(z)\Theta_{\text{LHS}}(z)]' + nf(z)F^{n-1}(z)x\\
&\qquad + n(n-1)(n-2)^2F^{n-3}(z)\big[\frac{1+F(z)}{2(n-2)} - \frac{F(z)}{n-1}\big][1-F(z)]f(z)x\\
&\qquad + n(n-1)(n-2)F^{n-2}(z)\times\\
&\qquad\quad \big[\frac{-2F(z)f(z)}{2(n-2)} - \frac{f(z)[1-F(z)]-F(z)f(z)}{n-1}\big]f(z)x\\
&\qquad + n(n-1)(n-2)zf(z)F^{n-2}(z)[1-F(z)]\\
&\qquad - n(n-1)(n-2)zf(z)F^{n-3}(z)\frac{1-F^2(z)}{2}\\
&\qquad - \frac{-2F(z)f(z)}{2}\int_{-\infty}^{z} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-3}(y_{n-2})dy_{n-2}\\
&\qquad - f(z)\int_{-\infty}^{z} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-2}(y_{n-2})dy_{n-2}\\
&\qquad + \frac{n!}{(n-3)!3!}F^{n-4}f(z)[1-F(z)]^3(x-z)\\
&\qquad - \frac{n!}{(n-3)!3!}F^{n-3}(z)3[1-F(z)]^2f(z)(x-z)\\
&\qquad - \frac{n!}{(n-3)!3!}F^{n-3}(z)[1-F(z)]^3\\
&\qquad + \frac{n!}{(n-4)!3!}F^{n-4}(z)[-1+3F(z)-3F^2(z)+F^3(z)]f(z)x\\
&\qquad + \frac{n!}{(n-4)!3!}F^{n-4}(z)[1-F(z)]^3f(z)z\\
&\quad = -[F_{Y_n}(z)\Theta_{\text{LHS}}(z)]' + nf(z)F^{n-1}(z)x\\
&\qquad - n(n-1)(n-2)F^{n-2}(z)\big[\frac{F(z)f(z)}{n-2} - \frac{F(z)f(z)}{n-1}\big]f(z)x\\
&\qquad + n(n-1)(n-2)F^{n-3}(z)\times\\
&\qquad\quad \big[\frac{1-F(z)^2}{2} - \frac{(n-2)F(z)[1-F(z)]}{n-1} - \frac{F(z)[1-F(z)]}{n-1}\big]f(z)x\\
&\qquad + n(n-1)(n-2)zf(z)F^{n-3}(z)[F(z)(1-F(z)) - \frac{1-F^2(z)}{2}]\\
&\qquad + \int_{-\infty}^{z} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-3}(y_{n-2})\times\\
&\qquad\quad [F(z)-F(y_{n-2})]dy_{n-2}f(z)\\
&\qquad + \frac{n!}{(n-3)!3!}F^{n-4}(z)[1-F(z)]^2\big[(n-3)(x-z)f(z) -\\
&\qquad\quad (n-3)(x-z)f(z)F(z) - 3(x-z)f(z)F(z) - F(z)(1-F(z))\big]\\
&\qquad - \frac{n!}{(n-4)!3!}F^{n-4}(z)[1-F(z)]^3f(z)x\\
&\qquad + \frac{n!}{(n-4)!3!}F^{n-4}(z)[1-F(z)]^3f(z)x].
\end{aligned}
$$

$$
\begin{aligned}
&\frac{\partial u(x,z|z\le x)}{\partial z}\\
&\quad= -[F_{Y_n}(z)\Theta_{\text{LHS}}(z)]' + nf(z)F^{n-1}(z)x - nf(z)F^{n-1}(z)x\\
&\qquad + n(n-1)(n-2)F^{n-3}(z)\frac{[1-F(z)]^2}{2}f(z)x\\
&\qquad - n(n-1)(n-2)F^{n-3}(z)\frac{[1-F(z)]^2}{2}f(z)z\\
&\qquad + \int_{-\infty}^{z} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-3}(y_{n-2})\times\\
&\qquad\quad [F(z)-F(y_{n-2})]dy_{n-2}f(z)\\
&\qquad + \frac{n!}{(n-3)!3!}F^{n-4}(z)[1-F(z)]^2\times\\
&\qquad\quad \big[(n-3)(x-z)f(z) - n(x-z)f(z)F(z) - F(z)(1-F(z))\big]\\
&\qquad + \frac{n!}{(n-4)!3!}F^{n-4}(z)[1-F(z)]^3 f(z)(z-x)\\
&\qquad \ge 0.
\end{aligned}
$$

At $z = x$, the above reduces to

$$
\begin{aligned}
&\frac{\partial u(x,z|z\le x)}{\partial z}|_{z=x}\\
&\quad= -[F_{Y_n}(x)\Theta_{\text{LHS}}(x)]'\\
&\qquad + \int_{-\infty}^{x} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-3}(y_{n-2})\times\\
&\qquad\quad [F(x)-F(y_{n-2})]dy_{n-2}f(x)\\
&\qquad - \frac{n!}{(n-3)!3!}F^{n-3}(x)[1-F(x)]^3\\
&\quad= 0.
\end{aligned}
$$

Thus, we have

$$
\begin{aligned}
&[F_{Y_n}(x)\Theta_{\text{LHS}}(x)]'\\
&\quad= \int_{-\infty}^{x} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-3}(y_{n-2})\times\\
&\qquad [F(z)-F(y_{n-2})]dy_{n-2}f(x)\\
&\qquad - \frac{n!}{(n-3)!3!}F^{n-3}(x)[1-F(x)]^3.
\end{aligned}
\tag{6.14}
$$

We have now obtained a closed form solution for computing both $\Theta_{\text{RHS}}(x)$ and $\Theta_{\text{LHS}}(x)$. Combining equations (6.13) and (6.14), and noting that $\frac{n!}{(n-3)!3!}F^{n-3}(x)[1-F(x)]^3 > 0$, we see that

$$
\begin{aligned}
[F_{Y_n}(x)\Theta_{\text{LHS}}(x)]' &= \int_{-\infty}^{x} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-3}(y_{n-2}) \times \\
&\quad [F(z)-F(y_{n-2})]dy_{n-2}f(x) \\
&\quad -\frac{n!}{(n-3)!3!}F^{n-3}(x)[1-F(x)]^3 \\
&< \int_{-\infty}^{x} n(n-1)(n-2)y_{n-2}f(y_{n-2})F^{n-3}(y_{n-2}) \times \\
&\quad [F(z)-F(y_{n-2})]dy_{n-2}f(x) \\
&= [F_{Y_n}(x)\Theta_{\text{RHS}}(x)]'.
\end{aligned}
$$

Because $\beta_{\text{RHS}}(0) = \beta_{\text{LHS}}(0) = \beta(0) = 0$, $F_{Y_n}(0)\Theta_{\text{RHS}}(0) = F_{Y_n}(0)\Theta_{\text{LHS}}(0) = 0$. Thus, integrating both side of the above equation, we obtain for all x

$$\Theta_{LHS}(x) < \Theta_{RHS}(x),$$

and thus for all x

$$\beta_{LHS}(x) < \beta_{RHS}(x),$$

which implies that β does not exist for any x. This result contradicts the assumption that β is monotonically increasing and continuous with $\beta_{LHS}(x) = \beta_{RHS}(x)$. Thus, in the symmetric, sequential Vickrey auctions with all bids revealed, there does not exist a symmetric pure-strategy equilibrium. $\diamond$

6.4 Non-Existence of Asymmetric Equilibrium

Having shown that symmetric equilibrium do not exist in our sequential auction model, we now turn our attention to the existence of asymmetric equilibria. In an *asymmetric equilibrium*, bidders are not restricted to using identical strategy functions. We present an example that shows that asymmetric equilibrium are not guaranteed to exist in the sequential FPSB auction model, while leaving unproven the question of whether they ever exist.

Lemma 6.4.1. *In the symmetric sequential FPSB auction model with bid revelation, if an asymmetric, pure-strategy equilibrium exists, it may not be unique.*

	Using $\{\beta_1, \beta_2, \beta_3\}$	Deviate to $z = 0$	Increase in Utility When Deviating
$x > Y_2$	U_1	$x - \epsilon$	> 0
$Y_2 > x > Y_1$	U_2	$x - \epsilon$	> 0
$Y_1 > x$	0	0	$= 0$
Overall			> 0; better off

Table 6.3: The overall utilities of bidder 0 in the sequential auctions.

Proof: The proof is trivial. Suppose there are three bidders and there exists an asymmetric equilibrium $\{\beta_1, \beta_2, \beta_3\}$. Given that these three strategy functions are different from one another, immediately, we have five other asymmetric equilibria: $\{\beta_1, \beta_3, \beta_2\}$, $\{\beta_2, \beta_1, \beta_3\}$, $\{\beta_2, \beta_3, \beta_1\}$, $\{\beta_3, \beta_1, \beta_2\}$, $\{\beta_3, \beta_2, \beta_1\}$. $\diamond$

Theorem 6.4.2. *The symmetric sequential auction model with bid revelation may not have any asymmetric pure-strategy equilibrium generically, when all bidders stick to the same asymmetric strategy profile throughout the game.*

Proof: We prove the theorem by contradiction. Assume that in the symmetric sequential auction model with bid revelation there exists an asymmetric pure-strategy equilibrium. Now consider the following special case. Suppose there are three bidders and there exists an asymmetric equilibrium $\{\beta_1, \beta_2, \beta_3\}$. We show that there exists a strategy that is better than $\{\beta_1, \beta_2, \beta_3\}$ or any of its variations implied by Lemma 6.4.1.

We again assume that the bidder under consideration is bidder 0, which has a true valuation x. Without loss of generality, we assume $\beta_1(x) < \beta_2(x) < \beta_3(x)$. This assumption allows these three strategy functions cross each other; however, these three strategy functions should be monotonically increasing. Similar to Example 6.3.1, we show that bidder 0 can be better off by unilaterally deviating from the asymmetric equilibrium $\{\beta_1, \beta_2, \beta_3\}$, as demonstrated in Table 6.3.

In Table 6.3, U_1 is the utility that bidder 0 gains in the whole game when $x > Y_2$. U_2 is the utility that bidder 0 gains in the whole game when $Y_2 > x > Y_1$. In detail,

$$U_{1,2} = \Pr(\text{wins 1})\big[x - \beta_1(x)\big] + \Pr(\text{loses 1; wins 2})\big[x - E[Y_{Lose}|\text{loses 1; wins 2}] - \epsilon\big],$$

where Y_{Lose} is the highest valuation of the other bidders who lose in the first auction. Because $\Pr(\text{wins 1}) + \Pr(\text{loses 1; wins 2}) < 1$, we may show that $U_1 < x - \epsilon$ as long as $E[Y_{Lose}|\text{loses 1; wins 2}]$ is positive and ϵ goes to infinitesimal. Similarly, $U_2 < x - \epsilon$ as long as $E[Y_{Lose}|\text{loses 1; wins 2}]$ is positive and ϵ goes to infinitesimal. When $Y_1 > x$, bidder 0 will not be better off by unilaterally deviating.

Overall, bidder 0 will be better when unilaterally deviating to 0 in the first auction. This conflicts the definition of equilibrium that no bidder can be better off by unilaterally deviating. Thus, in our symmetric model, there does not generally exist an asymmetric equilibrium. $\diamond$

Since we use a special case to disprove the existence of an asymmetric equilibrium in the general symmetric model, it does not rule out that there exists an asymmetric equilibrium in some special case of this symmetric model. It does not rule out that there might exist an equilibrium that is neither symmetric equilibrium nor asymmetric equilibrium either.

6.5 Conclusions

In this chapter, we consider a variation of the classic symmetric, sequential-auction model in which all bids are revealed after each auction, a market structure that is quite common in public marketplaces such as eBay. We show that there does not exist a pure-strategy equilibrium in either first- or second-price auctions. We also discuss the non-existence of asymmetric equilibrium in the sequential first-price sealed bid auctions. These results do not rule out mixed-strategy equilibria. Although the majority of existing literature is focused on sequential auction models that do have equilibria, we show in this chapter that the existence of a pure-strategy equilibrium is not guaranteed in some important classes of sequential auctions.

We recognize that the assumption that strategy functions be continuous, monotone, strictly increasing, and invertible, though commonly used in sequential auction models, is quite restrictive. Relaxing the assumptions may lead to different results. Also, changing other assumption might also affect the outcomes. For example, if the second auction in a two-stage sequential auction were a Vickrey auction, it is easy to prove that there exists a symmetric equilibrium despite the reveal-all-bid price quote. This result holds because, in the second auction, bidders will bid at their true valuations whether they have complete or incomplete information about the other bidders' valuations. Similar solutions can be found in much of the literature on second-price sequential auctions, see for example [15].

Chapter 7

Summary and Future Work

7.1 Summary of Contributions

With the bolstering of e-business, auctions have played an even more important role in trading in both business-to-business and public marketplaces. Many auction mechanisms have been introduced to cater to the demands of commerce. Bidders, naturally, demand tools to aid their strategic decision making in these auction games. This thesis aims to provide strategies for buyers/bidders in single item and sequential auction models.

Finding closed-form solutions to some auction games is challenging. Sequential auctions introduce more computational complexity by adding multi-stages. Decision-making tools in this multi-agent, multi-stage environment requires not only economic, but computational efficiency as well.

The existing rich volume of literature on sequential auctions provides answers to a variety of scenarios. This thesis contributes to this trend. However, theoretic results often cannot be easily applied to general sequential auction games. Although we prefer closed-form solutions, when they are not forthcoming, we seek sound heuristic approaches. Thus, we aim to design a flexible decision-making system for solving a broad class of auction games, with either discrete or continuous bids.

My contributions include the analysis of some typical individual auctions. More specifically, I study the multiplicity of equilibria in first-price sealed-bid (FPSB) auctions with discrete bids and complete information. I show that there are at most three equilibria and at least one equilibrium in two-agent FPSB auctions. I also discuss the equilibrium in sequential FPSB auctions with discrete bids. Further, I provide solutions to the FPSB auctions with discrete bids and incomplete information. We expect that a different tie-breaking rule might play a role in the final results.

While there are still plenty of unsolved sequential auction games, it is natural to design computational tools to solve a broader number of sequential auction models. In this thesis, I presented a Monte Carlo simulation method in approximating solutions for a group of diverse sequential auction games. I show how Monte Carlo sampling can be deployed to construct a bidding policy that performs comparably to the subgame perfect equilibrium. This method takes advantage of information revealed in previous auctions in the sequence to improve play in later auctions. The leveraged structure of the extensive form game, as a representation of the sequential auctions with discrete bids, is used to save computation memory and computation time dramatically. For example, an original extensive form game with 4.5 billion decision nodes can be reduced to a leveraged structure with only 1931 decision nodes. Importantly, the architecture is flexible, in that it can comprise a variety of auction models, and different types of bidders. I also prove that this Monte Carlo approximation approach converges to the average policy of perfect Bayesian equilibrium.

Information naturally plays an important role in finding optimal solutions to auction games. In sequential auctions, information revealed in the previous rounds might be used to help decision-making for the next ones. This becomes obvious in my model of symmetric sequential auction, in which all bidders and their bids are revealed after each auction. Revealing bids after an auction is quite common in public marketplaces, such as eBay and other online auction sites. Although the majority of existing literature is focused on sequential auction models that have equilibria, I prove that there does not exist a pure-strategy symmetric equilibrium in both sequential, first-price, sealed-bid auctions and sequential Vickrey auctions. I also discuss the non-existence of pure-strategy asymmetric equilibrium in the symmetric first-price sealed-bid auctions. This work provides a road map for future study in sequential auctions with continuous bids when bids are revealed.

7.2 Future Work

The literature on sequential auctions often relies on some strict assumptions. In real marketplaces, we often find that the number of agents, the number of items for sale, and the order of the auctions are stochastic, rather than static. I plan to study the impact of these stochastic factors in finding strategic equilibria. Meanwhile, it would be interesting to study the impact of other parameters, such as the reserve price, budget constraints, buy-it-now features, and other options in sequential auctions.

While concluding the non-existence of pure-strategy, symmetric equilibrium when bids are revealed, the existence of equilibrium in the model remains unanswered. Since the strategy space is infinite and the information is incomplete, Nash's theorem cannot be applied to this model directly. It remains to be seen whether some mixed strategy equilibria exist for these classes of sequential auction games.

A sequential auctions with incomplete information is a special case of sequential games under uncertainty. It might be possible to extend the approaches in this thesis to sequential games with uncertainty. Meanwhile, partially observable Markov decision processes (POMDP) have been widely adopted for these types of problems. I intend to study whether we might use POMDP for solving sequential auctions.

Finally, the majority of research in sequential auctions is focused on finding closed-form solutions theoretically. As sequential auctions become more and more popular, the attention of researchers will shift to more empirical study. The interdisciplinary study between auctions and other areas, such as supply chain management, is also promising.

Appendices

Appendix A

Mathematics Prerequisites

Bayesian Rule

Theorem A.1 (Multiplication rule). *Given that A_1 and A_2 are events, we have*

$$P[A_1 \cap A_2] = P[A_2|A_1]P[A_1] = P[A_1|A_2]P[A_2]. \tag{A.1}$$

Theorem A.2 (Bayes' theorem [75]). *Given that A_i, $i = 1, ..., n$, are a collection of events which partition A, and B is an event such that $P[B] \neq 0$. Then, for any $j \in \{1, ..., n\}$, we have*

$$P[A_j|B] = \frac{P[B|A_j]P[A_j]}{\sum_{i=1}^{n} P[B|A_i]P[A_i]}. \tag{A.2}$$

Order Statistics

Consider order statistics $Y_1 \leq Y_2 \leq ...Y_r \leq Y_s \leq ... \leq Y_n$ from a same cumulated distribution function (CDF) F and the corresponding probability density function (pdf) f. The probability distribution function of Y_r is given by

$$f_{Y_r}(y_r) = n! f(y_r) \frac{[1 - F(y_r)]^{n-r}}{(n-r)!} \frac{[F(y_r)]^{r-1}}{(r-1)!}.$$

The joint pdf of Y_r and Y_s is given by

$$f_{Y_r,Y_s}(y_r, y_s) = n! f(y_r) f(y_s) \frac{[F(y_r)]^{r-1}}{(r-1)!} \frac{[F(y_s) - F(y_r)]^{s-r-1}}{(s-r-1)!} \frac{[1-F(y_s)]^{n-s}}{(n-s)!}.$$

Typically, we have

$$f_{Y_n}(y_n) = nF(y_n)^{n-1} f(y_n),$$

and its associated CDF is

$$F_{Y_n}(y_n) = F(y_n)^n.$$

The joint pdf of Y_n and Y_{n-1} is

$$f_{Y_{n-1},Y_n}(y_{n-1}, y_n) = n(n-1) f(y_n) f(y_{n-1}) F(y_{n-1})^{n-2}.$$

The joint pdf of Y_{n-1} and Y_{n-2} is given by

$$f_{Y_{n-2},Y_{n-1}}(y_{n-2}, y_{n-1}) = n(n-1)(n-2) F(y_{n-2})^{n-3} f(y_{n-2}) f(y_{n-1}) [1 - F(y_{n-1})].$$

Appendix B

Notation

A A set of all agents.

b_j^i a specific jth bid by agent i.

B_i Bidding space of agent i.

C A constant value.

$f(x)$ A function of variable x.

f_i Probability density function of agent i.

F_i Accumulated density function of agent i.

J A set of other agents.

h_k^i Bidding history of agent i at kth auction.

$H_k^{\mathcal{J}}$ Bidding history of other agents $\mathcal{J}$ at kth auction.

I Set of information.

$\mathcal{J}$ A subset of other agents.

l An instance of samples.

L A set of experiment samples.

$P_r(x)$ A probability conditional on x.

s A strategy.

x The value of a variable X.

X_i A variable for agent i.

y_i The concrete value of Y_i.

Y_i The ith order statistics.

u_i The utility function of agent i.

v_i A specific valuation of agent i.

β_i The bidding function of agent i.

γ A component game representation.

Γ A game representation.

ξ An information state.

Π A strategy set.

σ A strategy profile.

$\hat{\sigma}$ A mixed strategy profile.

Σ A set of strategy profiles.

ϕ_i The reversed bidding function of agent i.

$\omega_i^l(\gamma)$ The policy for agent i at subgame γ in instance l.

Ω_i^l The policy for agent i in instance l.

Bibliography

[1] P. Anthony, W. Hall, V. Dang, and N. R. Jennings. Autonomous agents for participating in multiple on-line auctions. In *IJCAI Workshop on E-Business and the Intelligent Web*, pages 54–64, Seattle, WA, 2001.

[2] Susan Athey. Single crossing properties and the existence of pure strategy equilibria in games of incomplete information. *Econometrica*, 69:861–890, 2001.

[3] Howard James Bampton. Solving imperfect information games using the monte carlo heuristic. Technical report, University of Tennessee, Knoxville, 1994.

[4] Alan Beggs and Kathryn Graddy. Declining values and the afternoon effect: Evidence from art auctions. *The RAND Journal of Economics*, 28:544–565, 1997.

[5] Dan Bernhardt and David Scoones. A note on sequential auctions. *The American economic review*, 84:653–657, 1994.

[6] Alexander Bogomolny. In some circumstances index equals the content, 2003. $http : //www.cut - the - knot.com/do_you_know/poincare.shtml$.

[7] Craig Boutilier, Moises Goldszmidt, and Bikash Sabata. Sequential auctions for the allocation of resources with complementarities. In *Sixteenth International Joint Conference on Artificial Intelligence*, pages 527–534, Stockholm, 1999.

[8] Fernando Branco. Sequential auctions with synergies: An example. *Economics Letters*, 97:159–163, 1997.

[9] Andrew Byde, Chris Preist, and Nick Jennings. Decision procedures for multiple auctions.

In *Proceedings of the first International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002)*, pages 613–620, 2002.

[10] Gangshu Cai and Peter Wurman. Monte Carlo approximation in incomplete-information, sequential-auction games. *Decision Support Systems*, 39:153–168, 2005.

[11] Sandra Campo, Isabelle Perrigne, and Quang Vuong. Asymmetry in first-price auctions with affiliated private values. *Journal of Applied Econometrics*, 18:179–207, 2003.

[12] S. Chung. Np-completeness of the linear complementarity problem. *Journal of Optimization Theory and Applications*, 69:393–399, 1989.

[13] Richard W. Cottle, Jong-Shi Pang, and Richard E. Stone. *The Linear Complementarity Problem*. Academic Press, 1992.

[14] B. C. Eaves. Homotopies for computation of fixed points. *Mathematical Programming*, 3:1–22, 1972.

[15] Wedad Elmaghraby. The importance of ordering in sequential auctions. *Management Science*, 49:673–682, 2003.

[16] Richard Engelbrecht-Wiggans. Optimal auctions revisited. *Games and Economic Behavior*, 5:227–239, 1993.

[17] Richard Engelbrecht-Wiggans. Sequential auctions of stochastically equivalent objects. *Economics Letter*, 44:87–90, 1994.

[18] Richard Engelbrecht-Wiggans, Paul R. Milgrom, and Robert J. Weber. Competitive bidding and proprietary information. *Journal of Mathematical Economics*, 11:161–169, 1983.

[19] Richard Engelbrecht-Wiggans and Robert J. Weber. An example of a multi-object auction game. *Management Science*, 25:1272–1277, 1979.

[20] Richard Engelbrecht-Wiggans and Robert J. Weber. A sequential auction involving asymmetrically-informed bidders. *International Journal of Game Theory*, 12:123–127, 1983.

[21] Ian Frank, David Basin, and Hitoshi Matsubara. Monte-carlo sampling in games with imperfect information: Empirical investigation and analysis. In *Game Tree Search Workshop*, 1997.

[22] Drew Fudenberg, D. Kreps, and D. Levine. On the robustness of equilibrium refinements. *Journal of Economic Theory*, 44:354–380, 1988.

[23] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1991.

[24] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, 1996.

[25] Ian L. Gale and Donald B. Hausch. Bottom-fishing and declining prices in sequential auctions. *Games and Economic Behavior*, 7:318–331, 1994.

[26] Ian L. Gale, Donald B. Hausch, and Mark Stegeman. Sequential procurement with subcontracting. *International Economic Review*, 41:989–1020, 2000.

[27] Ian L. Gale and Mark Stegeman. Sequential auctions of endogenously valued objects. *Games and Economic Behavior*, 36:74–103, 2001.

[28] GameTheory.net. Dictionary: Definitions of economic & game theory terminology, 2005. http://www.gametheory.net/Dictionary/.

[29] Amy Greenwald and Peter Stone. Autonomous bidding agents in the trading agent competition. *IEEE Internet Computing*, April 2001.

[30] Kemal Guler and Bin Zhang. Bidding by empirical bayesians in sealed bid first price auctions. Technical Report HPL-2002-212, HP Laboratories, Palo Alto, 2002.

[31] John C. Harsanyi. Games with incomplete information played by bayesian players. *Management Science*, 14:159–182,320–334,486–502, 1967-8.

[32] Donald B. Hausch. A model of sequentical auctions. *Economics Letters*, 26:227–233, 1988.

[33] M.D. Hirsch, C.H. Papadimitriou, and S.A. Vavasis. Exponential lower bounds for finding brouwer fixed points. *Journal of Complexity*, 5:379–416, 1989.

[34] Robert D. Hof. The eBay economy. *Business Week*, August 25, 2003.

[35] Shlomit Hon-Snir, Dov Monderer, and Aner Sela. A learning approach to auctions. *Journal of Economic Theory*, 82:65–88, 1998.

[36] Thomas D. Jeitschko. Learning in sequential auctions. *Southern Economic Journal*, 65:98–112, 1998.

[37] Thomas D. Jeitschko. Equilibrium price paths in sequential auctions with stochastic supply. *Economics Letters*, 64:67–72, 1999.

[38] Ehud Kalai and Dov Samet. Persistent equilibria in strategic games. *International Journal of Game Theory*, 1984.

[39] Brett Katzman. A two stage sequential auction with multi-unit demands. *Journal of Economic Theory*, 86:77–99, 1999.

[40] V. Klee and G.J. Minty. *How Good is the Simplex Algorithm?*, pages 159–175. Academic Press, New York, 1972.

[41] Paul Klemperer. Auction theory: Aguide to the literature. *Journal of Economics Surveys*, 13:227–286, 1999.

[42] Paul Klemperer. Auction theory: A guide to the literature. In Paul Klemperer, editor, *The Economic Theory of Auctions*, pages 3–62. Edward Elgar, 2000.

[43] Elon Kohlberg and Jean-Francois Mertens. On the strategic stability of equilibria. *Econometrica*, 54:1003–1038, 1986.

[44] Daphne Koller and Nimrod Megiddo. The complexity of two-person zero-sum games in extensive form. *Games and Economic Behavior*, 4:528–552, 1992.

[45] Daphne Koller, Nimrod Megiddo, and Bernhard von Stengel. Fast algorithms for finding randomized strategies in game trees. In *26th ACM Symposium on the Theory of Computing*, pages 750–759, 1994.

[46] Daphne Koller, Nimrod Megiddo, and Bernhard von Stengel. Efficient computation of equilibria for extensive two-person games. *Games and Economic Behavior*, 14(2):247–259, 1996.

[47] Daphne Koller and Avi Pfeffer. Representations and solutions for game-theoretic problems. *Artificial Intelligence*, 94(1):167–251, 1997.

[48] David Kreps and Robert Wilson. Sequential equilibria. *Econometrica*, 50:863–894, 1982.

[49] David M. Kreps. *Nash Equilibrium*, pages 167–177. The Macmillan Press, 1989.

[50] Kala Krishna. Auctions with endogenous valuations: The persistence of monopoly revisisted. *The American Economic Review*, 83:147–160, 1993.

[51] Vijay Krishna. *Auction Theory*. Academic Press, 2002.

[52] Val E. Lambson and Norman K. Thurston. Sequential auctions: Theory and evidence from the seattle fur exchange. a working paper, 2003.

[53] Bernard Lebrun. Existence of an equilibrium in first price auctions. *Economic Theory*, 7:421–443, 1996.

[54] Bernard Lebrun. First price auctions in the asymmetric n bidder case. *International Economic Review*, 40:125–142, 1999.

[55] Bernard Lebrun. Continuity of the first price auction Nash equilibrium correspondence. *Economic Theory*, 20:435–453, 2002.

[56] C. E. Lemke and J. T. Howson. Equilibrium points of bimatrix games. *Journal of the Society of Industrial and Applied Mathematics*, 12:413–423, 1964.

[57] Dan Levin and James L. Smith. Equilibrium in auctions with entry. *The American Economic Review*, 84:585–599, 1994.

[58] Alessandro Lizzeri. Uniqueness and existence of equilibrium in auctions with a reserve price. *Games and Economics Behavior*, 30:83–114, 2000.

[59] Eric Maskin and John Riley. Asymmetric auctions. *Review of Economic Studies*, 67:413–438, 2000.

[60] Eric Maskin and John Riley. Equilibrium in sealed high bid auctions. *Review of Economic Studies*, 67:439–454, 2000.

[61] David McAdams. Characterizing equilibria in asymmetric first-price auctions. Mimeo, 2003.

[62] R. Preston McAfee and John McMillan. Auctions and bidding. *Journal of Economic Literature*, 25:699–738, 1987.

[63] R. Preston McAfee and Daniel Vincent. The declining price anomaly. *Journal of Economic Theory*, 60:191–212, 1993.

[64] R. Preston McAfee and Daniel Vincent. Sequentially optimal auctions. *Games and Economic Behavior*, 18:246–276, 1997.

[65] Richard McKelvey. A liapunov function for Nash equilibria. Technical report, 1991. California Institute of Technology.

[66] Richard McKelvey and Andrew McLennan. Computation of equilibria in finite games. In H. Amman, D. A. Kendrick, and J. Rust, editors, *The Handbook of Computational Economics*, volume I, pages 87–142. Elsevier Science B.V., 1996.

[67] Richard McKelvey, Andrew McLennan, and Theodore Turocy. Gambit – software tools for game theory. Technical report, 2003. http://econweb.tamu.edu/gambit/manual-0.97.1.0/gambit.html.

[68] Richard McKelvey and Tom Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10:6–38, 1995.

[69] Richard McKelvey and Tom Palfrey. Quantal response equilibria for extensive form games. *Experimental Economics*, 1:9–41, 1998.

[70] Andrew McLennan. The asymptotic expected number of nash equilibria of two player normal form games. a working paper, 2002.

[71] Andrew McLennan. The expected number of nash equilibria of a normal form game. a working paper, 2002.

[72] Paul R. Milgrom. Auctions and bidding: A primer. *Journal of Economic Perspectives*, 3:3–22, 1989.

[73] Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50:1089–1122, 1982.

[74] Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding II. In Paul Klemperer, editor, *The Economic Theory of Auctions*, volume 2, pages 179–194. Edward Elgar Publishing Limited, 2000.

[75] J.S. Milton and Jesse C. Arnold. *Introduction to Probability and Statistics: Principles and Applications for Engineering and the Computing Sciences*. McGraw-Hill, 1995.

[76] K. G. Murty. Computational complexity and linear pivot methods. *Mathematical Programming Study*, 7:61–73, 1978.

[77] Roger Myerson. Refinements of the Nash equilibrium concept. *International Journal of Game Theory*, 7:73–80, 1978.

[78] Roger Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1997.

[79] Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:59–73, 1981.

[80] Roger B. Myerson. Nash equilibrium and the history of economic theory. *Journal of Economic Literature*, 36:1067–1082, 1996.

[81] John Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36:48–49, 1950.

[82] John Nash. Two-person cooperative games. *Proceedings of the National Academy of Sciences*, 21:128–140, 1950.

[83] John Nash. Noncooperative games. *Annals of Mathematics*, 54:289–295, 1951.

[84] Elmor L. Peterson. *Zero-Sum Two-Payer Games (Matrix Games)*, 2000. Course Handout of Linear Programming.

[85] Carolyn Pitchik and Andrew Schotter. Perfect equilibria in budget-constrained sequential auctions: and experimental study. *The RANK Journal of Economics*, 19:363–388, 1988.

[86] Armando Ortega Reichert. A sequential game with information flow. In Paul Klemperer, editor, *The Economic Theory of Auctions*. Edward Elgar Publishing Limited, 2000. *Models for Competitive Bidding Under Uncertainty*, Chapter VIII, Ph.D. thesis and Technical Report No.8, January, 1968, Department of Operations Research, Stanford University, 232-54.

[87] John G. Riley and William F. Samuelson. Optimal auctions. *The American Economic Review*, 71:381–392, 1981.

[88] J. Rosenmuller. On a generalization of the Lemke-Howson algorithm to noncooperative n-person games. *SIAM Journal of Applied Mathematics*, 21:73–79, 1971.

[89] Alvin E. Roth and Axel Ockenfels. Last-minute bidding and the rules for ending second-price auctions: Evidence from ebay and amazon auctions on the internet. forthcoming.

[90] John Rust, John H. Miller, and Richard Palmer. Characterizing effective trading strategies: Insights from a computerized double auction tournament. *Journal of Economic Dynamics and Control*, 18:61–96, 1994.

[91] R. Saigal. On the convergence rate of algorithms for solving equations that are based on methods of complementary pivoting. *Mathematics of Operations Research*, 2:108–124, 1977.

[92] Herbert Scarf. The approximation of fixed points of a continuous mapping. *SIAM Journal on Applied Mathematics*, 15:1328–1343, 1967.

[93] Herbert Scarf. *The Computation of Economic Equilibria*. New Haven: Yale University Press, 1973.

[94] Reinhard Selten. Re-examination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory*, 4:25–55, 1975.

[95] Reinhard Selten. Evolutionary stability in extensive two-person games - correction and further development. *Mathematical Social Sciences*, 16:223–266, 1988.

[96] Harshit S. Shah, Neeraj R. Joshi, and Peter R. Wurman. Mining for bidding strategies on eBay. Submitted for publication, May 2002.

[97] Søren Tang Sørensen. Optimal sequential auctions for complements. Technical report, NERA, London, 2003.

[98] Peter Stone, Michael L. Littman, Satinder Singh, and Michael Kearns. ATTac-2000: An adaptive autonomous bidding agent. *Journal of Artificial Intelligence Research*, 15:186–206, 2001.

[99] A.J.J. Talman and Z. Yang. A simplicial algorithm for computing proper equilibria of finite games. Center discussion paper 9418, Tilburg University, 1994.

[100] Gerry Tesauro and Rajarshi Das. High-performance bidding agents for the continuous double auction. In *IJCAI Workshop on Economic Agents, Models, and Mechanisms*, 2001.

[101] G. van der Laan, A. J. J. Talman, and L. van Der Heyden. Simplicial variable dimension algorithms for solving the nonlinear complementarity problem on a product of unit simplices using a general labelling. *Mathematics of Operations Research*, pages 377–397, 1987.

[102] William Vickrey. Conterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

[103] Nils-Henrik Morch von der Fehr. Predatory bidding in sequential auctions. *Oxford Economic Papers*, 46:345–356, 1994.

[104] Bernhard von Stengel. Efficient computation of behavior strategies. *Games and Economic Behavior*, 14:220–246, 1996.

[105] Bernhard von Stengel. Computing equilibria for two-person games. In *Handbook of Game Theory with Economic Applications*, volume 3, pages 1723–1759. Elsevier, Amsterdam, 2002.

[106] Robert J. Weber. Multiple-object auctions. In R. Engelbrecht-Wiggans, M. Shubik, and R. M. Stark, editors, *Auctions, Bidding and Contracting: Uses and Theory*, pages 165–91. New York University Press, 1983.

[107] Michael P. Wellman, Amy Greenwald, Peter Stone, and Peter R. Wurman. The 2001 trading agent competition. In *Fourteenth Conference on Innovative Applications of Artificial Intelligence*, pages 935–941, Edmonton, 2002.

[108] Michael P. Wellman, Peter R. Wurman, Kevin A. O'Malley, Roshan Bangera, Shou-De Lin, Daniel Reeves, and William E. Walsh. Designing the market game for a trading agent competition. *IEEE Internet Computing*, 5(2):43–51, Mar.-Apr 2001.

[109] Peter R. Wurman. *Market Structure and Multidimensional Auction Design for Computational Economies*. PhD thesis, University of Michigan, 1999.

[110] Peter R. Wurman. Toward flexible trading agents. *AAAI Spring Symposium on Game Theoretic and Decision Theoretic Agents*, pages 134–140, March, 2001.

[111] Peter R. Wurman, William E. Walsh, and Michael P. Wellman. Flexible double auctions for electronic commerce: Theory and implementation. *Decision Support Systems*, 24:17–27, 1998.

[112] Peter R. Wurman, Michael P. Wellman, and William E. Walsh. A parametization of the auction design space. *Games and Economic Behavior*, 35:304–338, 2001.

[113] Zaifu Yang. *Computing Equilibria and Fixed Points*. Kluwer Academic Publishers, 1999.

[114] Weiying Zhang. *Game Theory and Information Economics*. Shanghai People's Press, 1996. A Chinese version.

[115] Weili Zhu and Peter R. Wurman. Structural leverage and fictitious play in sequential auctions. In *Eighteenth National Conference on Artificial Intelligence*, pages 385–390, 2002.