THE SECOND CHANCE OFFER: SELLER AND BIDDER STRATEGIES

by Yu-An Sun

B. S. in Computer Science, June 1999, National Tsing-Hua University, Taiwan

M.B.A. in Information Management, June 2001, National Central University, Taiwan

A Dissertation submitted to

The Faculty of
School of Engineering and Applied Science
of The George Washington University
in partial fulfillment of the requirements
for the degree of Doctor of Science

January 31st, 2009

Dissertation directed by

Dr. Poorvi L. Vora
Assistant Professor of Computer Science

The School of Engineering and Applied Science of The George Washington University certifies that Yu-An Sun has passed the Final Examination for the degree of Doctor of Science as of October, 9th, 2008. This is the final and approved form of the dissertation.

THE SECOND CHANCE OFFER: SELLER AND BIDDER STRATEGIES

Yu-An Sun

Dissertation Research Committee:

Dr. Poorvi L. Vora, Assistant Professor of Computer Science, Dissertation Director

Dr. Hyeong-Ah Choi, Professor of Engineering and Applied Science, Committee Member

Dr. Shmuel Rotenstreich, Associate Professor of Engineering and Applied Science, Committee Member

Dr. Rahul Simha, Professor of Engineering and Applied Science, Committee Member

Dr. Sumit Joshi, Professor of Economics, Committee Member

DEDICATION

To My Family,

Especially My Grandfather,

Mr. Nai-Shun Sun,

who never stops encouraging me to read more, learn more and think more;

who never loses his confidence in me at every step of the way.

None of my contributions would be possible without your support and love.

ACKNOWLEDGMENTS

The author wishes to acknowledge the support and guidance of her advisor, Dr. Poorvi L. Vora for telling her to do the work she likes and to do it well, Dr. Sumit Joshi of the Department of Economics for teaching her much of what she knows about game theory, and all the professors in the Department of Computer Science at the George Washington University. The author also wishes to thank all her friends for their long-lasting patience and support.

ABSTRACT

The second chance offer is a common seller practice on eBay. It consists of price discrimination against the losing bidder, who is offered an identical item at the value of his or her highest bid. Prior work has shown that, if the price discrimination is certain—that is, the items are always offered to bidders at their highest losing bids—bidders can predict it, and it results in revenue loss for the seller. This dissertation hence allows the seller to randomize his strategy. It examines a similar, more general problem: a seller has k items. They are sold to n bidders in a two-stage game. The first stage is a sealed-bid private-value auction with n bidders. The second stage is a take-it-or-leave-it offer to each of $k-1$ losing bidders; randomized between a fixed-price offer and a second-chance offer. Showing that analytic techniques do not provide complete solutions because bidding strategies are not always monotonic increasing, this dissertation uses genetic algorithm simulations to determine the Bayesian (near-Nash) equilibrium strategies for bidders and sellers, for $n = 8$ and different values of k. It analyzes item scarcity and two types of auction mechanisms for the first stage: first-price auction and second-price auction. It tests the approach on real eBay data, and a rational bidding tool is implemented to illustrate the practical use of this model on eBay. This dissertation's use of randomized seller strategies and genetic algorithm simulations is unique in the study of the second-chance offer.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1: INTRODUCTION

It is no longer novel to exchange personal information for merchandise or money. For example, millions of consumers trade personal information-- such as home address, age, gender --for grocery discount membership. Popular e-commerce sites, such as Amazon, offer a discount coupon feature, the "gold box" for items related to an individual's previous purchase, encouraging individuals to allow their purchases to be linked. At first glance, and to the naïve consumer, the revelation of personal information does seem to only provide a financial benefit; however, immediate gratification may have a negative impact in the later stage. For example, filling out a questionnaire online in exchange for a free t-shirt will lead to numerous junk mails in the future. A study has shown that consumers tend to not take later consequences into consideration [Acquisti 04]; it is possible that this is because future impact is too ambiguous to take into immediate consideration. Thus, privacy, especially in the numerous possible online interactions, is not well understood, either by researchers or by consumers. Researchers do not have a formalization of the privacy problem that balances the benefits with the costs, and, as a result, consumers do not completely understand the consequences of revealing information, nor, if they did understand the consequence, do they have a means of determining a best response.

This dissertation presents a quantitative privacy model based on game theory. The model takes into account the fact that information revelation may provide advantages, as well as bear costs. In the model, privacy is treated as protection from information revelation in a multi-stage game. Interactions online, with merchants or other service providers, may be viewed as stages of a multi-stage game. Actions in one stage typically reveal information about the player, and result in positive or negative repercussions in a later stage. The dissertation applies this model to a special type of online interaction: eBay's second-chance offer.

Note that, if future impact is not ambiguous, as it is now, rational agents can help players, such as consumers, to play rationally. One may hence imagine the model presented in this dissertation as enabling a privacy infrastructure where, in addition to security tools for the protection of personal information, rational agent tools helped a consumer decide how much information to reveal, when, and to whom. Today, however, the basic building blocks for such an infrastructure do not exist. For example, how would one derive a strategy given a set of stages (auctions, retail sales, etc.) in a game? How does one determine when it is useful to protect a player in a game (like a bidder in an auction) and not reveal the identity of the player? When is it useful for the player to be recognized? Can we build automatic tools with optimal bidding strategies in mind? This dissertation focuses on a generalization of eBay's second-chance offer to make more specific the privacy problem and to provide strategies for the consumer that best balance benefit with cost.

1.1 The Second-Chance Offer and a Generalization

Consider an open-cry first-price private-value auction such as on eBay. If the auction is a stand-alone game not connected to any other, and timing effects are ignored, the best strategy for each bidder is to bid higher than other bidders, up to x, his valuation. Thus, at the end of the auction, only the winning bidder has not bid as high as his valuation. On the other hand, the highest losing bids of all other bidders, who have dropped out of the bidding, reflect their valuations. This strategy may be exploited by the seller through a second stage of price discrimination where a losing bidder is offered the item at his highest failed bid. This is commonly used on eBay, where it is termed the *second-chance offer*. When bidders do not anticipate the second-chance offer, the highest failed bids are bidder valuations, and the second stage charges the highest possible acceptable price to the losing bidder. On the other hand, if the losing bidder were to obtain the item through another auction held by the seller, the bidder would pay a lower price, as, in that auction, he would not bid as high as his valuation.

Several variations of the second-chance offer have been studied. In particular, Salmon and Wilson [Salmon and Wilson 06] study the problem when the second stage consists of an offer that inverts the known symmetric bidder strategy (including one that anticipates the second-chance offer). They find that the only existing Nash equilibria for two bidders are mixed-strategy equilibria. Joshi et al [Joshi, Sun and Vora 05] find that, when the number of bidders is greater than twice the number of available items, equilibria exist such that bidders lower their bids in anticipation. These observations lead naturally to the question of whether the seller can improve his revenue by randomizing the second-chance offer – that is, by not being predictable enough for the bidder to lower his bid substantially. This paper addresses a game, similar to the second-chance offer but more general, where the seller randomizes the second-chance.

A two-stage game is played between n bidders and one seller. In the first stage, all bidders enter a sealed-bid, private-value, first-price auction, and one item is sold. After the first stage is over, all non-winning bidders enter the second stage for $k-1$ items identical to the one sold in the first stage. In this stage, the seller makes a second chance offer with probability α, or a fixed-priced offer, P, with probability $1-\alpha$. Bidders accept any offer that is not greater than valuation. The dissertation demonstrates that analytic techniques, which assume monotonic increasing bidding strategies, do not provide complete solutions to the game. That is, standard analytic techniques do not provide the seller and bidder strategies in equilibrium. Without the assumption of monotonic increasing strategies, it is not possible to present a simple expression for the bidder's optimization criterion, and the problem is essentially one of trying all possible bidder strategies for all possible valuations for the other bidders, to determine the optimal bidding strategy in the game. Genetic algorithms are hence used to solve the two-stage game model and determine the equilibrium strategies for both bidders and sellers. Two-population genetic algorithm (GA) experiments are conducted where bidders and sellers form the two populations. When the population converges, we have obtained near Nash equilibria.

The results of the GA experiments demonstrate that when the first-price auction is adopted in the first stage, and items are scarce —that is, less than or equal to half of the number of bidders ($k \leq 4$ when $n = 8$)—optimal bidding strategies are monotonic increasing. Otherwise, the strategies are not monotonic increasing. Further, item scarcity motivates bidders to bid high, and price discrimination ($\alpha = 1$) is an optimal seller strategy when items are scarce. When the second-price auction is adopted in the first stage, the results indicate that it is optimal for sellers to always price discriminate, regardless of item scarcity.

In the following sections, we describe the model and the contributions of this dissertation in more detail.

1.2 The Model

There are *n* bidders and one seller in the two-stage game. In stage I, all bidders enter an auction in which the bidder with the highest bid wins the object at the value of the highest bid. In stage II, the seller makes a second chance offer to all non-winning bidders from stage I. Each bidder *i* has a different valuation, x_i, which is the highest she is willing to pay; we assume each x_i is independently and identically uniformly distributed in [0,1]. The goal for all bidders is to maximize the expected value of their payoff, which is $x_i - b_i$ for the winning bidder, and zero for all other bidders. The seller's goal is to maximize the expected value of its revenue, which is the value of the winning bid.

Bidders' valuations may be viewed as their private information. Bids in stage I reveal the valuations, and this information is later used against bidders, resulting in a lower payoff. Therefore, the second-chance offer is a form of *privacy-infringement*. Rational bidders change their bidding strategy in response to such privacy infringement, by bidding low in stage I. *[Joshi, Sun and Vora 05]*, which is not part of this dissertation, shows that the seller does not benefit from a deterministic strategy if the bidders know it beforehand, as bidders will bid very low in order to make up for the payoff loss of the second-chance offer. Thus there is no incentive for the seller to price discriminate with certainty, and a simple rational strategy provides privacy protection to the bidder in such a case. In fact, *[Joshi, Sun and Vora 05]* shows that rational behavior provides very strong protection of the payoff, and is preferable to cryptographic protection. It shows that rational bidding by intelligent agents without cryptographic protection yields the highest bidder payoff and the lowest total revenue for sellers – in comparison to mechanisms where cryptographic protection is used, without rational behavior. The intelligent bidding agent also allows bidders with low valuation to signal their valuations to the seller, which reduces the opportunity loss for sellers' total revenue.

1.3 Contributions and Findings

While closed-form solutions exist for several single-stage auctions – among them first and second-price sealed-bid auctions – the addition of a second stage that is dependent on the first one makes the problem more difficult. This work is the first to examine:

- a randomized seller strategy in an auction, and its impact on bidder privacy
- privacy games without explicit closed-form solutions.

The contributions of this dissertation are as follows.

- It proposes a quantitative model of privacy based on a game-theoretic approach, and applies the model to the specific problem of the second-chance offer on eBay.
- It obtains near-Nash equilibrium results for a general game similar to the second-chance offer, where the seller's strategy is randomized. In particular, it obtains solutions when the first stage is a first or second price auction.

- It presents an approach for obtaining solutions when standard Bayesian analysis does not provide a solution, and when standard assumptions on the monotonic-increasing nature of the bidding strategies do not hold. It allows us to obtain pure as well as mixed bidding strategies in auctions when analytical approaches do not provide solutions.
- This dissertation applies its results to real data obtained from eBay to demonstrate the efficacy of the proposed techniques. The results with real data also demonstrate the use of rationality as privacy protection.
- An automatic bidding tool that determines optimal bidding strategy has also been developed.

The results of this dissertation indicate that whether randomized or deterministic strategies are optimal for the seller depends on the auction mechanism used, and, when the first-price auction is used, also on the scarcity of items (that is, on the relationship of *k* and *n)*. When the first stage of the game is a first-price auction, we have found that bidders do not penalize a privacy-infringing seller as much when fewer items are available and valuations are distributed as is standard in theoretical auction models, uniformly in [0, 1]. This is because competition among bidders for fewer items motivates the bidder to bid higher. Thus, when items are scarce, rational behavior does not provide sufficient privacy protection. However, when we examine the case of real valuation distributions estimated from eBay data, as well as uniformly distributed between the highest and lowest valuations estimated from eBay data, with the second-price auction in stage I, the results indicate that it is always optimal for sellers to price discriminate—that is, use the privacy infringing option in stage II. This result holds regardless of item scarcity.

This dissertation is organized as follows: Chapter 2 contains related work, Chapter 3 presents our model. Chapter 4 presents the Bayesian analysis, and Chapter 5 presents genetic algorithm simulations assuming bidder valuations are distributed uniformly. Chapter 6 presents results using real eBay data to determine the valuation distribution, and the bidding tool implementation. Chapter 7 contains conclusions and directions for future research.

CHAPTER 2: RELATED WORK

There are four major research areas that are related to our work. First, our work views privacy as a game of information revelation; hence the first area related to our work studies how individuals view their privacy, and how dynamic pricing affects their decisions to reveal information, because price discrimination can be viewed as a form of privacy invasion. This research area explores the economic aspects of personal privacy. We describe it in further detail in section 2.1.

Second, our work models privacy as information revelation in a two-stage game that includes the first-price or second-price auctions. It is a game of incomplete information, and players strategically signal their valuation to maximize payoff. We describe research on relevant game-theoretic aspects of auctions and games of incomplete information in section 2.2. Third, our work looks at the specific privacy game of the second chance offer. We describe work related to the second-chance offer in sections 2.3 and 2.4. In section 2.3, we discuss a deterministic model that is similar to our two-stage game. The only difference is that, in the deterministic model, sellers' action to price discriminate in certain. Our two-stage game can be viewed as a generalization of the deterministic model. In section 2.4, we describe research on a randomized strategy for privacy. These two sections explore various models that can be considered as variations of our model, which generalizes them. We describe various well-known optimization methods in section 2.5, which are typically used to obtain optimal strategies.

Finally, when examining optimal strategies in our two-stage game of incomplete information, we adopt the evolutionary programming method, also called genetic algorithms, and conduct experiments to obtain equilibria. The fourth related research area, of using genetic algorithms to solve economic problems, as well as some other applications of genetic algorithms, are described in section 2.6.

2.1 Personal Privacy and Dynamic Pricing

There is a great deal of work related to the economic aspects of personal privacy, for example, [Laudon 96] [Acquisti, Dingledine and Syverson 03] [Ackerman, Cranor and Reagle 99] [Acquisti and Grossklags 04]. Varian [Varian 96] was perhaps the first to propose that privacy be treated as an individual's right to property, where personal data forms property. In this framework, private information can be traded, sold or exchanged through market mechanisms. Varian also pointed out that once a consumer's private information is sold to a third party, no control is left for the original parties of the transaction. Varian concluded that to complete the framework of treating privacy as property rights, it is necessary to have legislations to regulate the secondary usage of privacy; it must have the consent of the original party.

Acquisti applied psychology and behavioral economics to analyze whether consumers make rational decisions regarding their privacy [Acquisti 04]. Acquisti introduces the idea of "immediate gratification"– consumers have self-control problems and lean towards obtaining immediate compensation without considering the long-term effect of their decisions. Acquisti also noted that the "rational privacy" model, in which the agent is assumed to have rationality and unbounded computational power, does not hold because of such psychological distortion. Those consumers who claim to value their privacy do not demonstrate this in their actions. In conclusion, Acquisti recognized the need to develop software tools, policies and government regulations to help consumers to make rational decisions regarding their privacy.

In addition, Acquisti and Varian examined whether it is profitable for a firm to perform first-degree price discrimination, i.e. condition prices based on purchase history [Acquisti and Varian 01]. In their analyses, it is shown that first-degree price discrimination is only profitable when there are a large number of uninformed consumers or when a firm can provide additional services for different value consumers. Acquisti and Varian classified consumers into two types: those with low values, and those with high values. They point out that price discrimination is achievable if the seller's additions are such that either consumers do not switch types, or only low value types switch to high value types.

On the topic of dynamic pricing, Odlyzko pointed out that the seller will try to extract additional revenue if it is aware of the consumer's willingness to pay more for the same goods [Odlyzko 03]. This happened in the railroad industry in the 19th century, as well as on e-commerce site Amazon.com in the 21st century. However, price discrimination requires a delicate balance because consumers resent obvious forms of price discrimination, while manufactures prefer it to maximize profit. Odlyzko noted that the requirement for balance results in a form of mild, stealthy price discrimination, such as product bundles, with the help of tools such as DRM (Digital Rights Management) to reduce consumer resentment as much as possible.

Odlyzko's work points out an important problem: when privacy protection only benefits consumers, vendors do not have any incentives to provide such protection. On the other hand, if privacy protection also benefits the seller, the seller will be motivated to provide the protection.

2.2 Auctions and Strategic Signaling

An auction is a widely-used pricing mechanism used for the allocation of goods [Klemperer 04] [Menezes and Monteiro 05]. McAfee and McMillan have summarized real world examples, auction theory development and types of auctions used [McAfee and McMillan 87]. There are four types of auctions: the English auction, the Dutch auction, the first-price sealed-bid auction and the second-price sealed-bid auction.

In an English auction, buyers continuously raise the price until there is only one buyer left. The winner pays the highest current bid. The English auction is also known as the *open-cry* auction because the current bid is always revealed to all participants. In a Dutch auction, the seller announces the initial price and keeps lowering the price until one buyer accepts it. In a first-price sealed-bid auction, all the bidders simultaneously submit their sealed bids to the seller; the highest bidder wins the item and pays its submitted bid. Similarly, in a second-price sealed-bid auction, all the bidders simultaneously submit their sealed beds; the highest bidder obtains the item, but pays the second-highest submitted bid.

The second-price sealed-bid auction is also known as the Vickrey auction, and was first theoretically analyzed by William Vickrey in 1961. In his attempt to theoretically analyze market mechanisms, Vickrey proposed an auction mechanism to sell the item to the highest bidder at the value of the second highest bid [Vickrey 61]. Vickrey auctions result in some nice properties – for example, the weakly dominant strategy is to bid one's true valuation, because one's bid only determines whether one loses or wins this auction. This yields an expected revenue (for the seller) that is equivalent to that of the English auction, and the auction is strategically equivalent to the Dutch auction.

Despite all the nice theoretical properties of Vickrey auctions, Vickrey auctions are used only rarely [Rothkopf, Teisberg and Kahn 90] [Lucking-Reiley 00]. Rothkopf observed that this is because the bidders fear that, if bid-takers use the information revealed by the bids in a future interaction, this could be of disadvantage to the bidder. This observation has inspired our work to model information revelation as a multi-stage game, because it is realistic for bidders to encounter one another again in practice. Dominant strategies derived from one-shot game assumptions do not reflect this valid concern.

McAfee and McMillan noted that a number of variant forms of these four auctions are commonly used in the real world, such as while imposing a "reserved price", charging an entry fee, and offering limited time to submit bids, etc. McAfee and McMillan further proved that these four auctions yield identical expected revenue for the seller, i.e. they proved the "revenue equivalence theorem". McAfee and McMillan also included the U.S. Treasury bonds auctions as real world examples. Discriminatory auctions and uniform-priced auctions are both used; however, discriminatory auctions are used to sell relatively shorter-term bonds while uniform-priced auctions are used to sell long-term bonds. It is also important to note that U.S. Treasury bond auctions are common-value auctions, meaning information of the bond value is publicly available.

In another paper, Milgrom compared the four different types of auctions–the English auction, the Dutch auction, the first-price sealed bid auction and the second-price sealed bid auction–and concluded that the English auction is popular because of the low participating costs and the sealed-bid auction has the risk of the seller inserting fake bids to raise the final price [Milgrom 89]. . Based on Milgrom's paper, various cryptographic auction schemes with a third party as the auctioneer are proposed.

In addition to auctions, strategic signaling is also relevant to our work. Crawford and Sobel proposed one of the first strategic communication models of the sender and receiver game [Crawford and Sobel 82]. The sender has private information, *m*, which is drawn by nature and not known to the receiver. The sender sends a costless, non-verifiable message to the receiver. The receiver acts according to the message and its action determines the payoff for both sender and receiver. This model is also called "cheap talk" because the communication is costless. Crawford and Sobel also showed that there is no equilibrium that does not contain a noisy message unless the sender and receiver have identical utility functions. In other words, revealing the whole truth is not optimal for the sender if its interest differs from that of the receiver. Crawford and Sobel then characterized another solution set as "partition equlibria", where the optimal strategy is to send a noisy message.

The problem we consider– that of finding a balance between consumer privacy and seller revenue – is similar to a generalized sender-receiver game, where more than two players are involved, with different interests in mind.

2.3 The Deterministic Model

In [Joshi, Sun and Vora 05], the deterministic version of the two-stage game very similar to that described in chapter 3 is studied. In Stage II, the seller either chooses the price-discrimination offer option with probability one, or has another auction to sell the other items. For a seller with identical objects; there are different mechanisms, including the second chance offer, that offer various degrees of privacy protection to the bidder. To compare the impact of these mechanisms on the bidder, the authors define the privacy cost as the payoff difference between different mechanisms for identical items. They present the payoff differences among various cases as depicted in Table 2-1:

Table 2-1: Payoff Difference Comparison

Case Zero	This is the baseline case where all the objects are sold in consecutive, independent auctions. The privacy costs of Cases B-D are defined wrt this case.
Case A	Bids and corresponding identities are all known to the seller. After the auction, the seller provides a second chance offer to all the bidders that didn't win the object. Bidders are naïve, and bid as though there is no second chance offer. This is expected to roughly correspond to current bidding on eBay.
Case B	The same as case A except bidders are strategic. This corresponds to bidder behavior if rational agents were available.
Case C	Only the bids are known to the seller, not the corresponding identities. The seller can only contact bidders as a group (and hence with a fixed price offer).
Case D	The seller only knows the final highest price and can only contact bidders as a group with a fixed price offer.

Cases B, C and D represent different kinds of assistance that may be provided a bidder:

- Case B corresponds to the use of intelligent bidding agents.
- Case C corresponds to the use of anonymity technology, which destroys any link between a bidder and its bid. Though individual bids are known, the

bidder corresponding to the bid is anonymous in the set of all bidders participating in the auction.

- Case D corresponds to the use of both: anonymity and bid secrecy technology, which does not provide information on losing bids to the seller.

For the seller, Cases C and D correspond to different levels of information on individual bidders and bids:

- In Case C, the seller can optimize its fixed price offer with the knowledge of all the bids.
- In Case D, the seller can only estimate a best fixed price offer based on an assumption of, say, uniformly distributed valuations, and make an offer equal to the midpoint of the highest bid.

The authors found that the seller's total revenue in case C is the same or larger than that in case D. Also, case A generates the highest revenue for the seller; and, interestingly, case B generates the lowest revenue and highest bidder payoff. That is, *the use of intelligent agents is more beneficial for bidder privacy protection than the use of cryptography*. The authors also found that this difference is pronounced for low valuation bidders. This is because intelligent agents allow the bidder to signal a low valuation to the seller, while cryptographic technology does not provide the option of doing so, and creates an opportunity loss when low valuation bidders are not able to signal their inability to buy at an average fixed price.

The above findings are not the contribution of this dissertation; these are included as related work because of their relevance to this dissertation: a generalization of the above game, to the two-stage game with probabilistic actions in stage II. While these findings demonstrate that it is not beneficial for the seller to price discriminate with certainty if the bidder is rational, this dissertation examines the consequences of the randomization of seller strategy.

2.4 Randomized Strategies for Privacy

In a paper that follows the first use of randomization in the second-chance offer, a contribution of this dissertation, [Joshi, Sun and Vora 08] examined a multiple-buyer game with two stages. In stage one, all buyers submit a sealed-bid in response to a declaration of pricing rules by the seller. In stage two, the seller makes a take-it-or-leave-it offer to the buyer with the largest signal; this value need not be equal to the signal. The authors show that, if the seller breaks the rules and price discriminates with certainty, the buyers reveal no information, but that, if the seller breaks the rules with a probability smaller than one, buyers reveal information in signals that increase seller revenue. This work does not, however, correspond directly to eBay's second chance offer, where seller's are only allowed to charge a bid in the second stage, and not any other value.

Salmon and Wilson studied a similar game, the English-Ultimatum game, where the first stage consists of a first price auction, and the second stage consists of a take-it-or-leave-it offer [Salmon and Wilson 06]. The authors showed that there is only a mixed-strategy equilibrium in a 2-item-2-bidder, English-Ultimatum game. The differences between their model and ours are as follows:

1. In the English-Ultimatum game, the seller always chooses to price discriminate in the second stage, while our model allows sellers to choose between price discrimination and uniform-price offers.
2. In the second stage, the English-Ultimatum game allows the seller to make an offer that is different from the failed bids obtained from stage one, while our model does not allow such a change.
3. Salmon and Wilson examines a mixed strategy for bidders while we examine the randomized strategy for sellers.
4. Our model is for any number of bidders.

The sender-and-receiver game studied first by Crawford and Sobel also addresses informed and uninformed players [Crawford and Sobel 82]. The sender is the informed player and the receiver is the uninformed player. The game is played as follows: nature chooses the sender's type first, the sender chooses a message and the receiver chooses an action afterwards. The payoff of both players is affected by the actions/messages they chose. This is different from the problem we address, because, in our problem, bidders are the uninformed players and move first. Further, in Crawford and Sobel game, utility function of the sender is non-zero if the receiver takes an action that is greater than the sender's secret. In our problem, the utility function is zero if the receiver (seller) takes an action (makes an offer) that is greater than the sender's secret (bidder valuation).

2.5 Optimization Methods

In this section we discuss some standard optimization techniques that are typically used to solve games as well as for other applications.

2.5.1 Deterministic Line Search

Steepest descent is an old optimization technique that can be applied to continuous and differentiable functions. It was first proposed by Cauchy in 1874. It is an iterative process to choose a starting point and a neighbor point that the function decreases most quickly based on the first derivative [Chong and Zak 96]. Newton-Raphson method is similar to steepest descent because it also chooses a starting point and the neighbor point. The difference is that the choosing Newton-Raphson method chooses the next point based on both the first and second derivatives [Chong and Zak 96]. Line search methods have similar flow: (1) initialize a starting point, (2) determine a direction, (3) compute the distance to move towards that direction, (4) move to the new target point and check whether it reaches optimum, (5) repeat from (1) if it's not optimum.

Comparing the two methods, it can be observed that steepest descent generally has a more rapid convergence in the beginning because it uses first-order derivative, but Newton-Raphson method has a more rapid convergence at the end of the process. A hybrid method can adopt both by implementing the steepest descent at the start and finishing up the optimization process with the Newton.

2.5.2 Linear, Nonlinear and Quadratic Programming

Linear, nonlinear and quadratic programming methods are all used to solve constrained optimization problems. In all 3 methods, there is an objective function and a set of constraints. The goal is to find a solution that maximizes or minimizes the objective function and also satisfies the constraints. Linear programming has an objective function and a set of constraints that are both linear. Similarly, quadratic programming method has an objective function that is quadratic (squared variables) and a set of linear constraints. Nonlinear programming is used to solve a nonlinear objective function and a set of constraints. If the constraints are linear, it is called "linearly constrained optimization".

The simplex method is commonly used to solve linear programming problems. It requires the set of constraints to be written in matrix form [Chong and Zak 96]. It finds a feasible solution by computing the inverse matrix and gradually moves from a feasible solution to an optimal solution by finding an adjacent solution with better value (higher if maximizing, lower if minimizing) of the objective function.

In a nonlinear programming problem, if the objective function is convex and the set of constraints is a convex set, convex optimization method can be used. One of the convex optimization methods is based on the Lagrange multiplier theorem. Instead of matrix operations, it uses the derivatives of the objective function [Chong and Zak 96]. The quadratic programming problem can also be solved by convex optimization method.

2.5.3 Random Search

Simulated Annealing was first proposed by Metropolis et al [Metropolis 53] in 1953. It is a method combining deterministic local search and probabilistic moves to reach a global optimum; at each point, the process makes a probabilistic decision of whether to stay at the current point in the deterministic search algorithm, or to move to another point. This prevents deterministic search algorithms, such as steepest descent, from getting ``stuck'' in local optima. Genetic algorithms are also categorized as optimization techniques based on random search, and mimic how nature is believed to reach equilibria. We describe the genetic algorithm in detail in the next section.

2.6 The Genetic Algorithm

The genetic algorithm was first introduced by Holland and Goldberg [Holland 75] [Holland 92] [Goldberg 88]. It is based on Darwinian natural evolution theory: the growth of animals is mainly controlled by their genes, inherited from their parents. Instead of reproducing the same genes from one single source, the genes are actually a mix of those of both parents, with possible random changes, known as mutation.

Adopting the biological model of evolution, solutions in genetic algorithms are coded as *chromosomes.* Similar to the notion of the survival of the fittest in Darwinian theory, a *fitness function* determines survival in genetic algorithms. This could be an objective function or a subjective function defined by human decisions. Genetic algorithms are iterative algorithms, and possible solutions are iteratively selected or rejected based on their fitness. Beginning with a first set of possible solutions, pairs of accepted solutions generate solutions for the next iteration, using *crossover* and *mutation*. This process is repeated until good enough chromosomes are created, or time runs out.

Selection is crucial in the genetic algorithm because it decides how to obtain more copies from the better solutions. There are several different ways of performing selection: roulette wheel selection, tournament selection and truncation selection. Chromosomes with higher fitness scores will have a higher percentage in the roulette wheel selection. In the tournament selection, the algorithm will randomly pick any two (or more) chromosomes, compare fitness scores, and keep the best one. The truncation method is the most trivial one; it doubles the better half population and truncates the other half.

The cross over function mimics the biological reproduction process. It combines bits from good parents generated by the previous selection function. There are two methods for cross over: one point cross over and two points cross over. For a one-point crossover, the algorithm first randomly picks a point in the bit strings. All bits before that point are from one parent and all bits after that point are from the other parent. Two point cross over works similarly.

Mutation provides variation that is needed in the genetic algorithm to prevent it from being limited by its first (randomly chosen) set of solutions. One bit in the chromosome bit string is randomly flipped. Mutation can also be viewed as a random walk away from the original chromosome. Because mutation causes variation, global optimization can be achieved.

There are other algorithms adopted from the power of nature, such as observing how ants alter their path to build the shortest path finding algorithm [Dorigo 92].

In our model, the genetic algorithm scheme is used to find Bayesian-Nash equilibria among sellers and bidders in a multi-stage game of incomplete information. From an evolutionary game theory point of view, the survived strategies result in mutual best responses for all the players because the strategies with lower utility are eliminated during the evolving process; therefore, the Bayesian-Nash equilibria obtained from our experiments are also evolutionary-stable.

2.7 Applications of Genetic Algorithms to Problems in Economics

There have been several applications of genetic algorithms to the solution of economic problems. One application is mechanism design and evaluation. Cliff adopted the genetic algorithm approach to investigate an optimal mechanism for an online auction trading environment [Cliff 03] [Cliff 06] [Walia, Byde and Cliff 03]. Cliff discovered a hybrid auction market evolving in his experiments, and noted that software agents can be used in online auctions, and that current human-developed auctions are not necessarily optimal. . The hybrid auction market that evolved in his experiments is much more market efficient, and results in more overall market profit than any human-designed auction mechanisms. Similarly, Byde adopted genetic algorithms to evaluate various auctions including first-price and second-price sealed-bid auctions [Byde 03], and demonstrated the evolution of a hybrid mechanism. It is important to note that Byde established that the GA-based solution is optimal regardless of whether it is a human-trading market or an agent-based trading market. Byde also noted that the advantage of the genetic-algorithm-based approach is that not-theoretically-analyzable factors can be taken into consideration during the simulation process through the use of evolution. Similar to our approach, Byde used a 1-to-1 mapping table to represent the bidding strategy where the entry of the table is the bidders' valuation, and the outcome of the lookup table is the bid.

Genetic algorithms have also been applied to well-established problems in economics, such as the prisoner's dilemma, auctions, the cobweb model and other microeconomic problems [Riechmann 01] [Dawid 96] [Dawid 99]. In the prisoner's dilemma, defection is the dominant strategy if the game is only played once. Alexrod studied the iterated prisoner's dilemma (IPD) game with genetic algorithms [Alexrod 87]. In his experiments, each player has a chromosome consisting of three previous moves: a player chooses either cooperation or defection and there are four possible outcomes for each move. It is shown that the optimal strategies that evolved from the experiments have similar properties as TIT FOR TAT [Alexrod 84], a strategy submitted by Anatol Rapoport in a previous IPD strategy contest. Axelrod concluded that the genetic algorithm is an effective optimization technique in a large problem space.

Andreoni and Miller conducted genetic algorithm experiments to explain the anomaly of human auctions [Andreoni and Miller 95]. The bidders gradually learn the optimal strategy by evolving. Andreoni and Miller examined the evolved strategies in common value auctions, affiliated private value auctions and independent private value auctions over a period of 1,000 generations. Each generation consists of twenty rounds of auctions and the fitness function is defined as the total profit over twenty rounds. Andreoni and Miller also examined experiments with settings of 8 bidders and 4 bidders. Andreoni and Miller concluded that bidders in the experiments did reach Nash equilibrium in the common value, first price auctions. Andreoni and Miller noted it is relatively difficult to converge to equilibrium in auctions because of poor feedback in the auction environment.

The cobweb model is mostly used to describe the supply and demand equilibrium in agricultural market. Different from other markets, it takes a significant period of time for crops to grow; therefore farmers need to estimate the quantity they need to plant based on their forecast of the market price [Pindyck and Rubinfeld 04]. The cobweb theorem states that the market price will converge to the intersection of the supply-demand curves, the equilibrium, after a long period of time. Genetic algorithms have been adopted to conduct several simulations for different cobweb designs [Arifovic 94] [Franke 98] [Dwaid and Kopel 98]. Arifovic showed that the genetic algorithm can be used as a decision making and learning tool to achieve equilibrium price. Arifovic also noted that a GA-based approach does not require the agents to be intelligent to begin with; instead, an agent can keep updating its prior beliefs during the process to produce optimal solutions.

Dawid and Kopel also adopted a genetic algorithm approach to study two cobweb models. In one model the farmer (or firm) decides whether to stay in the market or exit before deciding the production quantity; while in the other model, the firm can only make decisions about production quantity [Dawid and Kopel 98]. Dawid and Kopel conducted simulations with different coding schemes and different designs of fitness functions, and discovered, surprisingly, a different result for each coding scheme. While such a finding indicates that the results cannot be generalized, Dawid and Kopel noted that the genetic algorithm approach is still useful to analyze economic problems because it always initiates a heterogeneous population to begin further simulations. A heterogeneous population represents asymmetric strategies. This is a major advantage over theoretical analysis because it allows us to examine more complex problems with less constraints and assumptions, for example, we do not need to assume symmetric strategies for all players in a game.

In summary, we described four related work areas: economic aspects of personal privacy and their relationship with price discrimination; the independent auction literature as a one-shot game and strategic signaling in another multi-stage game; different models that can be viewed as variation of ours; and well-known optimization methods, as well as applications of evolutionary programming, a method we use. All four areas are closely related to different aspects of our model described in Chapter 3.

CHAPTER 3: THE MODEL

We model a two stage price discrimination game—an auction stage followed by a seller-offering stage—as follows.

Game Price Discrimination

– Stage I: *N* bidders join a sealed-bid auction. All bidders simultaneously and independently make bids. The bidder with the highest bid wins the auction. On the occurrence of a tie, the winner is chosen at random. The remaining *N-1* bidders enter Stage II.

– Stage II: The seller offers an identical item to all the remaining bidders:

(1) The *failed bid (privacy-infringing) option:* With probability α, the price is the bidder's highest bid in Stage I.

(2) The *uniform price(privacy-protecting) option* With probability $1-\alpha$, the price is a uniform price *P* for all bidders. The bidders can reject or accept either offer.

Each bidder's payoff is calculated as the difference between the price paid for the item and the bidder's valuation. In this two-stage game, bidders are seeking to maximize their payoff and sellers are seeking to maximize their total revenue over stages. There are two different auction mechanisms used in Stage I: first-price sealed-bid and second-price sealed-bid, these are described in detail in section 3.2.

3.1 Notation

We follow the notation of Krishna [Krishna 02] and denote the private valuation by *x*, the bid by *b*, the payoff by Π, and the expectation operator by *E[.]*. β(x) is the optimal bidding function. *G(x)* denotes the probability that a given valuation *x* is the highest among *n* bidders; *g(x)* denotes its derivative. *H(x)* denotes the probability that a given valuation *x* is among the highest *k* ones, but is not the highest; *h(x)* denotes its derivative. *P* denotes the uniform-price offer made in Stage II, and *R* the revenue. x_i and b_i denote the i^{th} highest valuations and bids respectively.

3.2 Stage I

Stage I may be a first or second price sealed-bid auction.

3.2.1 First-Price Sealed Bid Auction

In Stage I, if a first-price sealed bid auction is adopted, the winner pays its bid. A bidder's expected payoff is written as

$$E[\prod] = G(x)(x-b) + H(x)[\alpha(x-b) + (1-\alpha)(x-P)]$$

3.2.2 Second-Price Sealed Bid Auction

In Stage I, if a second-price sealed-bid auction is used, the winner pays the second highest bid. A bidder's expected payoff is written as

$$E[\prod] = G(x)(x - \frac{\int_0^x bg(y)dy}{G(x)}) + H(x)[\alpha(x-b) + (1-\alpha)(x-P)]$$

where $G(x)$ denotes the probability that a given valuation x is the highest valuation among n bidders, $\frac{\int_0^x bg(y)dy}{G(x)}$ denotes the expected value of the second highest bidder's bid, and $H(x)$ denotes the probability that a given valuation x is among the highest k valuations, but not the highest. If the number of available items is less than the number of remaining bidders and the seller chooses to price discriminate, only the highest $k-1$ bidders receive an offer, where $k-1$ is the number of available items. If the seller chooses otherwise, the seller provides the uniform-price offer to all $n-1$ bidders, bidders then notify the seller whether the offer is rejected or accepted. The seller randomly selects $k-1$ bidders among those who accept the uniform-price offer.

3.3 Assumptions

We make the following assumptions:

1. x is independent and identically distributed with a uniform distribution over interval $[0, 1]$, with cumulative distribution function F, and the corresponding probability distribution function f. Hence,

$$G(x) = [F(x)]^{n-1}$$

and

$$H(x) = \sum_{i=1}^{k=1} C_i^{n-1} [F(x)]^{n-i-1} [1 - F(x)]^i$$

2. $\beta(x)$ is monotonic-increasing and identical for all bidders
3. Bidders are risk-neutral.

4. Bidders reject any offer exceeding valuation

3.4 Impact of Stages on Revenue and Payoff

3.4.1 Independent Second-Price Auction

In our model, the second-price auction can be adopted at the first stage in a two-stage game. The optimal strategy of the bidder in an independent second-price auction is well-known to be a "truth-revealing" strategy. It is a dominant strategy for the bidder to always bid its valuation; that is, the bidder cannot make a better payoff with another bidding strategy, independent of the strategies of other bidders. We briefly describe the reasons for this strategy.

Suppose, instead of submitting a bid x, the bidder submits bid b such that:

1. $b > x$:
   a. If the highest bid among other bidders, b', is such that $b' > b$, both bids of b and x result in the same payoff: zero, as the bidder does not win the auction with either strategy.
   b. If $b' < x$, both b and x are the highest bids, the bidder wins with either bid, and makes the same payoff $x-b'$ independent of the value of b.
   c. If $b > b' > x$ the bidder wins the auction in a situation where he would not have won with bid x. However, in this case, the bidder cannot afford the item, priced at b', and makes a zero payoff.
2. $b < x$:

a. If the highest bid among other bidders, *b'*, is such that *b'*> *x*, both bids of *b* and *x* result in the same payoff: zero, as the bidder does not win the auction with either strategy.
b. If *b'*< *b*, both *b* and *x* are the highest bids, the bidder wins with either bid, and makes the same payoff *x-b'* independent of the value of *b*.
c. If *x* > *b'*> *b* the bidder loses the auction with bid *b* and makes zero payoff. With bid *x*, the auction would have been won, with a non-zero payoff *x-b'*

The objective function for the independent second-price auction is x-b_2 for the bidder who wins, and zero for all other bidders. In an independent second-price auction, the seller's action is fixed *a priori*, and the seller does not have an optimization criterion.

3.4.2 Independent First-Price Auction

In our model, the first-price auction can also be adopted at the first stage in a two-stage game. The optimal strategy of the bidder will depend on the nature of the second stage, and will, in general, be different from that of the first-price auction played as an independent game. For the purposes of comparison, and to determine if the seller benefits from the second stage, we present a brief overview of the well-known results about first-price auctions played as independent games.

The first-price auction when played as an independent game does not posses a dominant strategy; its symmetric Nash equilibrium bidding strategy is well known to be the expected value of the second-highest valuation, conditional to *x* being the highest one:

$$\beta(x) = E[x_2 \mid x = x_1] = \frac{\int_0^x yG'(y)dy}{G(x)}$$

This optimal strategy is obtained by differentiating the expected payoff with respect to the bid, and then assuming all bidding strategies are identical (for a symmetric equilibrium) and then setting the derivative to zero to determine an extreme value of the payoff. Finally, it is shown that, if all other bidders bid according to this strategy, and a single bidder deviates, the deviating bidder obtains a smaller payoff than it would if it were to follow the strategy.

When *f(x)* is the uniform distribution,

$$\beta(x) = \frac{n-1}{n}x$$

The expected payoff is:

$$E[\prod(x)] = \int_0^x G(y)dy = \frac{F(x)^n}{n}$$

and the expected revenue:

$$E[R(x)] = E[x_2] = \frac{n-1}{n+1}$$

The objective function for the bidder in an independent first-price auction is $G(\beta^{-1}(b))(x - \beta(x))$.

3.4.3 Two Stage Model

We further analyze the objective functions for both the seller and the bidder in a two stage game. The seller's goal is always to optimize its total revenue; if Stage I is the first-price sealed-bid auction, the seller's objective function can be written as:

$$R(\alpha, P, \beta, x_1, x_2, \ldots x_n) = b_1 + \sum_{i=2}^{k} \alpha b_i + \sum_{x>P}^{k} (1-\alpha)P$$

where $b_i = \beta(x_i)$. If the second-price sealed-bid auction is adopted for Stage I, the seller's objective function can be written as:

$$R(\alpha, P, \beta, x_1, x_2, \ldots x_n) = b_2 + \sum_{i=2}^{k} \alpha b_i + \sum_{x>P}^{k} (1-\alpha)P$$

For the bidder, if the first-price sealed-bid auction is adopted in Stage I, payoff can be written as:

$$\Pi(\alpha, P, \beta, x_1, x_2, \ldots x_n) = \begin{cases} x - b_1 & \text{if wins at 1st stage} \\ x - [\alpha b + (1-\alpha)P] & \text{if wins at 2nd stage} \\ 0 & \text{otherwise} \end{cases}$$

If the second-price sealed-bid auction is adopted at Stage I, bidder payoff can be written as:

$$\Pi(\alpha, P, \beta, x_1, x_2, \ldots x_n) = \begin{cases} x - b_2 & \text{if wins at 1st stage} \\ x - [\alpha b + (1-\alpha)P] & \text{if wins at 2nd stage} \\ 0 & \text{otherwise} \end{cases}$$

To obtain the average payoff, we can sum up the payoff function over all possible combinations of x_i, b_i. In summary, the bidder's objective function can be written as:

$$E[\Pi(\alpha, P, \alpha_{est}, P_{est}, x_1, x_2, ...x_n)] = \sum_{all\ \ possible\ (x_1, x_2, ...x_n)} Q(\alpha, P, \alpha_{est}, P_{est}, x_1, x_2, ...x_n)$$

And the seller's objective function can be written as:

$$E[R(\alpha, P, \alpha_{est}, P_{est}, x_1, x_2, ...x_n)] = \sum_{all\ \ possible (x_1, x_2, ...x_n)} R(\alpha, P, \alpha_{est}, P_{est}, x_1, x_2, ...x_n)$$

To find the Nash equilibrium, we have to optimize the bidder's objective function $\sum_{all\ \ possible\ (x_1, x_2, ...x_n)} Q(\alpha, P, \beta, x_1, x_2, ...x_n)$ given (α, P) because the bidder does not have knowledge of the seller's move (α, P) while submitting the bid. We also have to optimize the objective function $\sum_{all\ \ possible (x_1, x_2, ...x_n)} R(\alpha, P, \beta, x_1, x_2, ...x_n)$ given β for the seller because the bids are submitted based on the bidder's estimate of the seller's move.

CHAPTER 4: BAYESIAN ANALYSIS

In this chapter we present the Bayesian analysis performed by us to obtain the optimal bidding strategy in some cases. We also present the results of genetic programming simulations to solve for the Bayesian Nash equilibria. Before we provide our results, we first review the independent first-price auction.

4.1 Bayesian Nash Equilibrium—First-Price Auction at Stage I

This dissertation examines the following two cases when the first-price auction is adopted at the first stage:

Case I: α and P represent the reputation of the seller and are known to the bidder. k identical items are available.

Case II: α and P form the seller's second-mover strategy and are unknown to the bidder. k identical items are available

4.1.1 Case I

Consider a simplified version of the game, where the bidder knows the sellers' values of the probability α and the uniform price P prior to the start of the auction. This may be thought of as a steady state setting in a repeated game, where the seller has chosen optimal values of α and P, and the bidder has learnt them over repeated interactions. α and P represent the reputation of the seller, and we may consider them as representing the distribution on the type of the seller. Just as the distribution on the bidder's valuation is known to the seller, so also the parameter α is known to the bidder, along with P. The bidders submit bids in the first stage, and, after the auction ends, the system flips a coin, biased according to the value of α, to choose between the uniform price option and the failed bid option. The result of the coin flip may be viewed as the type of the seller – much as the valuation x represents the type of the bidder – and is similarly unknown, *a priori*, to the other players.

Theorem 1: There is no dominant deterministic strategy for all bidders in the Game Price Discrimination when the first-price auction is adopted at the first stage.

Proof: Consider the strategy of bidding zero. It is a strongly dominant strategy (that is, its payoff is strictly greater than that of any other strategy) for non-highest-bid bidders who do not win Stage I, as the payoff is strictly greater than any other if the seller price discriminates in Stage II, and the payoff is the same as any non-zero bid if the seller offers a uniform-price instead. However, for the highest bidder, if all other bids are small enough (what is small enough depends on α and P), a greater payoff is obtained by bidding higher than all other bids and winning Stage I rather than bidding zero and risking paying P in Stage II. This is not as good as bidding zero for the non-highest bidders. Therefore, there is no dominant deterministic strategy for the bidder.

Assuming the bidder uses a symmetric Bayesian Nash equilibrium strategy, following Krishna's notation [Krishna 02], for bidders with valuation higher than P, the expected payoff for bid b and valuation x is as follows:

$$E[\prod] = G(\beta^{-1}(b))(x-b) + H(\beta^{-1}(b))[\alpha(x-b) + (1-\alpha)(x-P)]$$

$$\text{where } G(x) = x^{n-1} \text{ and } H(x) = \sum_{i=1}^{k=1} C_i^{n-1} x^{n-i-1}(1-x)^i$$

H(x) is the probability for valuation *x* to be among the *k* highest ones but not the highest one. Otherwise, if valuation *x* is less than fixed price *P*, the expected payoff is

$$E[\prod] = G(\beta^{-1}(b))(x-b) + H(\beta^{-1}(b))[\alpha(x-b)]$$

Differentiating both equations wrt *b*, and equating to zero gives the following optimal bidding strategies as equation (1). See appendix A for detailed steps and a proof that shows it is a symmetric Nash equilibrium strategy.

Theorem 2:

$$\beta(x) = \begin{cases} \dfrac{\int_0^x yG'(y)dy + \alpha\int_0^x yH'(y)dy + (1-\alpha)\int_P^x (y-P)H'(y)dy}{G(x)+\alpha H(x)} & x > P \\[2ex] \dfrac{\int_0^x yG'(y)dy + \alpha\int_0^x yH'(y)dy}{G(x)+\alpha H(x)} & else \end{cases} \quad (1)$$

is a Bayesian Nash equilibrium for Game Price Discrimination if G(x) + αH(x) ≠0, β(x)> 0 and monotonic increasing. Proof: See appendix A.

However, *β(x)* is not always monotonic increasing for all *α* and *P*. When *β(x)* is not monotonic increasing, it contradicts assumption (2) from section 3.3 and is not an invertible strategy. Therefore, determining the Bayesian strategy by differentiating the expected payoff function does not provide a solution for $\beta(x)$. We cannot use it to compute the revenue for each value of α and P to determine the optimal α and P for the seller. We hence need an alternate approach to determine the equilibrium seller strategy, and the corresponding bidder strategy. In particular, we need an approach that does not require the bidding strategy to be monotonic increasing.

4.1.2 Case II

In a less constrained situation, the two-stage game would be played sequentially, with the seller moving later than the bidders. That is, the seller would determine the values of α and P after receiving all the submitted bids. Further, again, the bidding strategy need not be monotonic increasing.

As closed form expressions for the payoff for arbitrary (non-monotonic increasing) strategies do not exist, we use genetic algorithm simulations to examine arbitrary bidding strategies. Based on the analysis of objective functions in section 3.4.3, we can see that the two objective functions for bidders and sellers depend on each other. The optimization methods we studied in chapter 2 including deterministic search, linear programming and quadratic programming do not provide a good fit for such problems. The next class of optimization techniques is heuristic search method including neural network, simulated annealing and genetic algorithms. We choose genetic algorithms and map bidders and sellers to two populations that evolve together. The bidding strategy is not assumed to be monotonic increasing. It is assumed that $0 \leq \beta(x) \leq x$. Because β is not necessarily monotonic increasing, it is possible that a bidder with a lower bid might be able to pay the fixed-price for the object while a bidder with a higher bid might not. Hence the fixed-price offer in Stage II, made with probability $1-\alpha$, is made not only to the $k-1$ highest bidders, but to all bidders. A random $k-1$ are chosen from all bidders who accept the fixed price offer. The solution space of player strategies is expected to be large, as the bidding strategy is not constrained to a particular form.

When bidding strategies are not monotonic increasing, there is no straightforward formula for the probability of winning, as it is not straightforward to characterize the distribution of the other bids; hence, for example, the probability of winning with bid b is not $G(\beta^{-1}(b))$ as in section 3.4.2 This further implies that the only formula for expressing the expected payoff is to provide an average over all possible valuations of the other bidders:

$$E[\prod(\beta(x), x)] = \begin{cases} \displaystyle\sum_{\beta(x)>\beta(x_i),\, x_i\in[0,1]} (x-\beta(x)) + \sum_{\beta(x)<\max\beta(x_i);\beta(x)\ among\ k-1\ highest\ bids, x_i\in[0,1]} \alpha(x-\beta(x)) + \\ (1-\alpha)(x-P)\Pr[bidder\ chosen\ for\ fixed\ price\ offer] & x>P \\ \displaystyle\sum_{\beta(x)>\beta(x_i),\, x_i\in[0,1]} (x-\beta(x)) + \sum_{\beta(x)<\max\beta(x_i);\beta(x)\ among\ k-1\ highest\ bids, x_i\in[0,1]} \alpha(x-\beta(x)) & else \end{cases}$$

We use genetic algorithm (GA) experiments to perform a search over the large solution space. In the GA experiments, the bidder's fitness function is payoff, and the seller's fitness function is revenue. The bidder's chromosomes define its strategy β, unconstrained except $0 \le \beta(x) \le x$. Using the GA experiments, the optimal strategies, for bidder and seller, are obtained when the chromosomes converge. Our approach and the experimental results are described in next section and chapter 5.

4.2 Evolutionary Programming to Determine Bayesian Nash Equilibrium

We first examine the case of no scarcity ($n = k$) where bidders follow the Bayesian strategy in equation (1). In this case, the number of bidders is the same as the number of available items.

4.2.1 Case I: $n = k$

First, for the purposes of illustration, consider Case I: when bidders known the values of α and P and $n = k$. Bidders' Bayesian strategy becomes equation (2) as in this case, $H(x) = 1-G(x)$, because $1-G(x)$ is the probability of that a bidder with valuation is among the highest k, $n = k$, bidders but not the highest one. It is a special case of (1). When $\alpha \neq 1$,

$$b = \beta(x) = \begin{cases} \dfrac{P}{1 + \dfrac{\alpha}{(1-\alpha)G(x)}} & \text{for } x \geq P \\ \\ \dfrac{\dfrac{n-1}{n}x}{1 + \dfrac{\alpha}{(1-\alpha)G(x)}} & \text{for } x < P \end{cases} \tag{2}$$

When $\alpha = 0$, bidders with valuations less than P can never obtain the item in the second stage. Low valuation bidders see the first stage as a simple first-price auction. On the other hand, bidders with valuations greater than P will never bid more than P in the first stage because they are guaranteed to be able to afford the price in Stage II. Given this response, it appears best for the seller to divide the bidders into two sets of roughly equal size, using *P=0.5* (as valuations are uniformly distributed), and to never price discriminate. This is confirmed in figure 4-1.

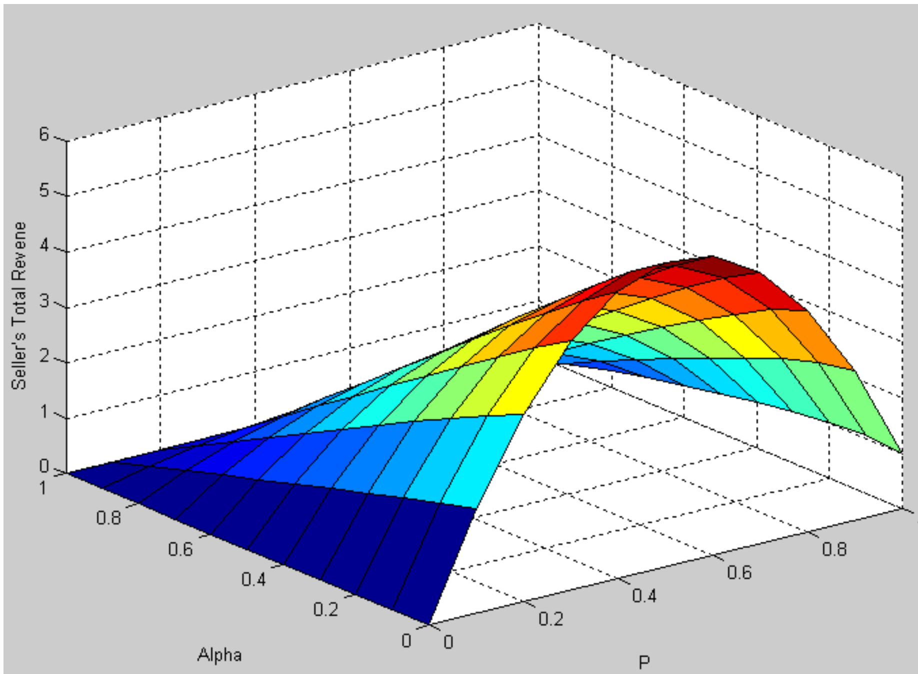

Figure 4-1: Seller's revenue as a function of α and *P*

Figure 4-1 shows the expected revenue computed from equation (2) as a function of α and *P*. The computation assumes uniformly distributed valuations, $n = 20$ bidders, and averages over 10 instances. It is clear that $\alpha = 0$ and $P = 0.5$ provide the optimal seller reputation, that is, the reputation at which the seller obtains the highest expected revenue. Note that $\alpha = 0$ and $P = 0.5$ provides a higher revenue than $\alpha = 0$ and $P = 1$, which provides the revenue of a simple first-price auction. This difference in revenue is largely due to the fact that $P = 0.5$ sells more items.

4.2.2 Case II: $n = k$

We now consider the case when the bidders do not know the values of α and P, but learn them in an evolutionary program, by estimating their values, and bidding according to the estimates. Good estimates will be passed on to the next generation. The sellers, simultaneously, through evolution, determine a best value of these parameters. In the experiment, the chromosomes represent the bidder's estimates of α and P for the fitness function *expected payoff*, and the seller's values of α and P for the fitness function *expected revenue*. The bidder's chromosome is coded as a pair of real numbers α and P. It represents the bidder's type, i.e. the bidder's assumed values of price discrimination probability α, and the uniform price P. The seller's chromosome is also coded as a pair of real numbers α and P. It represents the seller's action after Stage I ends.

We stop our simulation when the number of generations is 10,000. The initial population is set at 100 sellers, and 100 distinct bidders for each seller. Bidder chromosomes are randomly generated at the beginning of the simulation. Bidder valuations are randomly generated *each generation*. The tournament selection method is used for reproduction: two chromosomes will be randomly drawn from the population pool. The chromosome with the higher fitness score will be copied to the new population representing the next generation. It will stay in the pool for further tournaments, and the process continues until the new population has the same size as the previous one. The crossover process combines chromosomes from two parents, at random. The optimal mutation rate is set to be equal to

1/4 as suggested in [Mühlenbein and Schlierkamp-Voosen 93] because there are 4 different variables in the seller and bidder chromosomes. We implement the mutation operator suggested in [Mühlenbein and Schlierkamp-Voosen 93].

In the experiment, we divide the bidders into two groups: low valuation bidders and high valuation bidders. Group one contains bidders with low valuations, that is, valuations uniformly distributed between 0 and 0.5; group two contains bidders with high valuations, that is, valuations uniformly distributed between 0.5 and 1. The mutation rate is 1/4 from generation 0~5000 and 0.05*(1/4) afterwards. Bidders are assigned a randomly generated valuation every generation according to their group.

Table 4-1: Experiment Results

	Experiment: 100 Bidders	
	Mean	Variance
Sellers' α	0.0	0.0
Seller's *P*	0.5101	1.5660e-008
Low Valuation Bidder's α	0.4297	0.1188
Low Valuation Bidder's *P*	0.4434	0.1125
High Valuation Bidder's α	0.5866	0.1096
High Valuation Bidder's *P*	0.4155	0.1273

We find that the population of the seller's chromosome is converged. However, neither the low valuation bidders' nor the high valuation bidders' chromosomes are converged to any distinct pairs of (α, P). Table 4-1 summarizes the results of this experiment.

Figures 4-2 and 4-3 illustrate the bid distributions for high valuation bidders and low valuation bidders. The high valuation bidders' bids clearly converge to approximately zero. The low valuation bidders' population enters a stable stage since no significant change in population variance is found over generations. The only explanation for this result is that two Nash equilibria exist in this two stage game. One is to bid zero and one is to bid their valuation. The reasons why P and α do not converge are as follows:

- For high valuation bidders: although the bids converge to nearly 0, it appears that many combinations of α and P can result in nearly-zero bids.
- For low valuation bidders: because there are two Nash equilibria, α and P do not converge.

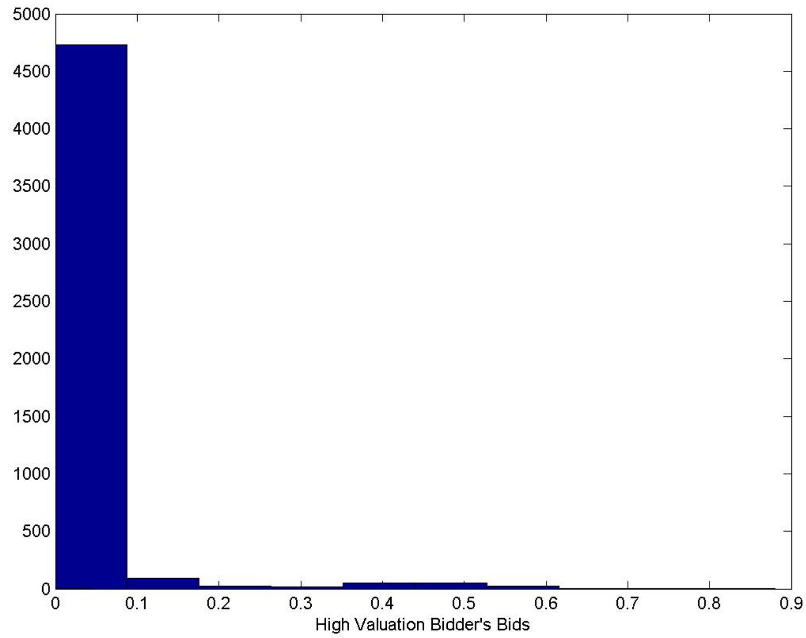

Figure 4-2: Distribution of High Valuation Bidder's Bids

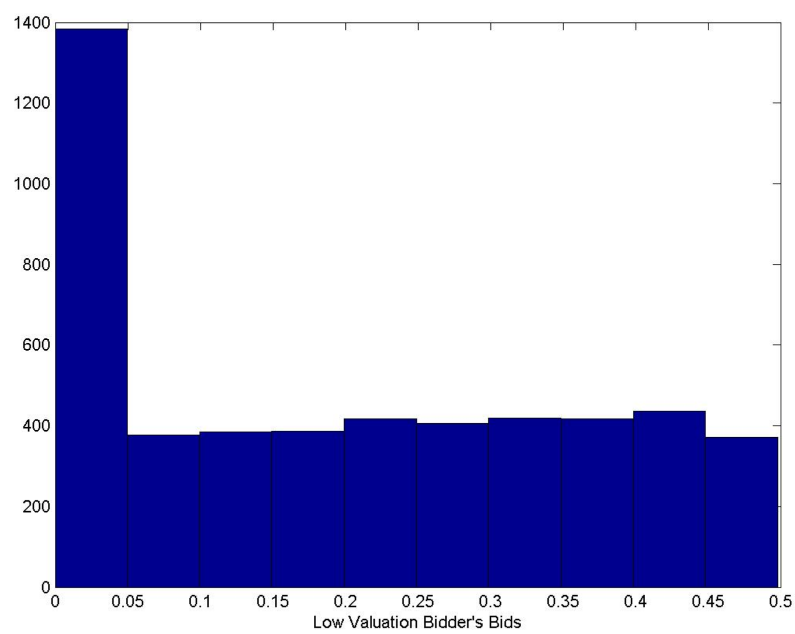

Figure 4-3. Distribution of Low Valuation Bidder's Bids

4.3 Bayesian Nash Equilibrium—Second-Price Auction at Stage I

This dissertation also examines the following two cases when second-price auction is adopted at the first stage: α and P represent the reputation of the seller and are known to the bidder. k identical items are available.

Similar to section 4.3, we consider a simplified version of the game, where the bidder knows the sellers' values of the probability α and the uniform price P prior to the start of the auction. The only difference is a second-price auction is adopted at Stage I.

Theorem 3: There is no dominant deterministic strategy for all bidders in the Game Price Discrimination when second-price auction is adopted at the first stage.

Proof: From section 3.4.1, we know that the dominant strategy to win a second-price auction is to bid up to one's true valuation. However, the dominant strategy for stage I results in the worst case at stage II when a second chance offer is made—with zero payoff. If the uniform-price is offered at stage II, the payoff is the same regardless of the bid. We can see that if a uniform-price offer is made at stage II, it is best to bid up to one's valuation at stage I to win the auction, if one's valuation is below P which depends on how the seller moves. However, bidding up to one's valuation does not result in an optimal outcome if a second-chance offer is made at stage II. Therefore, there is no dominant deterministic strategy exist for this game.

Assuming the bidder uses a symmetric Bayesian Nash equilibrium strategy, following Krishna's notation [Krishna 02], for bidders with valuation higher than *P*, the expected payoff for bid *b* and valuation *x* is as follows:

$$E[\prod] = G(\beta^{-1}(b))(x - \frac{\int_0^x bG'(y)dy}{G(x)}) + H(\beta^{-1}(b))[\alpha(x-b) + (1-\alpha)(x-P)]$$

$$\text{where } G(x) = x^{n-1} \text{ and } H(x) = \sum_{i=1}^{k=1} C_i^{n-1} x^{n-i-1}(1-x)^i$$

H(x) is the probability for valuation *x* to be among the *k* highest ones but not the highest one. Otherwise, if valuation *x* is less than fixed price *P*, the expected payoff is

$$E[\prod] = G(\beta^{-1}(b))(x - \frac{\int_0^x bG'(y)dy}{G(x)}) + H(\beta^{-1}(b))\alpha(x-b)$$

Differentiating both equations wrt *b*, and equating to zero results in partial differentiation equations from which we are not able to obtain an optimal strategy for general *k*. See appendix B for detailed steps and solutions for specific *n* and *k*.

Due to the difficulty of obtaining optimal strategies through solving partial differentiation equations, we conduct genetic algorithm experiments and the results are described in Chapter 5. It is worth noting that genetic algorithms have shown success in solving partial differentiation equations as mentioned in [Haupt and Haupt 04]. The genetic algorithm experiments include the case where this game is played sequentially, as well as the situation of item scarcity.

4.4 Summary

We have derived the conditions for injective non-negative strategies in Bayesian Nash equilibria for a two-stage price discrimination game when the first-price auction is adopted at stage I. In the absence of item scarcity, we find that rational behavior, even when the seller's strategy is randomized between price discrimination and fixed-price, provides benefit to the bidder and is sufficient for privacy protection, as it deters the seller from price discriminating.

In the next chapter, we examine item scarcity. We also examine the cases when the conditions under which we have determined the Bayesian Nash equilibria (injective, non-negative bidding strategies) are not obtained in the results; i.e. when the analytical approach does not obtain the equilibria, and there is no closed-form solution for it. In order to consider both cases, when the equilibria are characterized by closed-form solutions and when they are not, we characterize the bidders through their bidding strategies expressed as look-up tables, and not through their beliefs about α and P, (because even if they had correctly estimated α and P they would not have a closed-form bidding strategy).

CHAPTER 5: EVOLUTIONARY PROGRAMMING EXPERIMENTS

This chapter presents the detailed process and the results of the evolutionary programming experiments. Each experiment has various small values of n and k. Experiment set A adopted the first-price auction at stage I, and a uniform distribution of bidder valuations; for some values of n and k, results for this problem were also obtained through Bayesian analysis and are were described in Chapter 4. Due to the possible mixed strategies found in experiment A, we further conduct experiment B with most general chromosomal representation and uniform distribution to solve the game with the first-price auction in stage I. Finally, experiment set C adopted the second-price auction in Stage II (as would be the case on eBay), and used both uniformly distributed valuations, as well as valuations obtained from a real eBay dataset. Detailed results of experiment C are described in chapter 6.

5.1 Simulation Method

The fitness functions are expected payoff and expected revenue for bidder and seller respectively. The seller's chromosome consists of the values of α and P. In experiment set A, the bidder's chromosome is coded as a twenty-entry lookup table to represent a pure bidding strategy. The bidder's valuation is the index of the lookup table, and the bid is the content of the corresponding entry. In experiment set B, the bidder's chromosome is coded as two twenty-entry look up tables, and a probability γ for the first lookup table. The form of this chromosome is designed to represent a mixed strategy—with probability γ, the bid comes from the first lookup table, otherwise it comes from the second lookup table. In experiment set C, the bidder's chromosome is an *m* by *m* matrix. The $(i,j)^{th}$ element in the matrix represents the probability that the bid for valuation *i* will be the value *j*. Notice that the chromosome of experiment A is a special case of that of experiment C, and that of experiment C is a special case of that of experiment B.

5.2 Basic Steps

The stop condition is number of generations equals to 10,000 for experiment set A, and 5000 for experiment set B and C, which have finer chromosomal representations, and hence are expected to converge in fewer generations. The population is fixed across all experiments to 500 sellers, and $n = 8$ distinct bidders for each seller. Bidder chromosomes are randomly generated at the beginning of the simulation. Bidder valuations are randomly generated *each generation*. 1000 auctions are conducted during each generation for all three experiment sets. Because there is almost no literature on how to determine bidding functions that are not parameterized simple functional forms (such as linear or quadratic), we determined this number through several preliminary experiments, these are described in section 5.4.

The optimal mutation rate for experiment set A is set to be equal to 1/22 as suggested in [Mühlenbein and Schlierkamp-Voosen 93] because there are 22 different variables in the sellers and bidders chromosomes. We also implement the mutation operator suggested in [Mühlenbein and Schlierkamp-Voosen 93] as follows:

$$Var_i^{Mut} = Var_i + s_i \cdot a_i \cdot r_i, \, where \, i \in \{1,2,..10000\}$$

and

$$s_i \in \{-1, +1\}$$

$$r_i : mutation \quad range = 0.1*(1-0)$$

$$a_i = \sum_k u 2^{-k}, u \in \{0,1\} \; chosen \;\; at \; random \quad , k : mutation \;\; precision \; , \; k = 16$$

where i is the iteration number. A variable Var_i of the chromosome (each bidder chromosome contains 20 variables; each seller chromosome, two) is selected with probability *r=1/22* to mutate. For every variable, we flip a bias coin, with probability 1/22, the variable mutates[1]. For experiment B, the mutation rate is set to be 1/42 and $1/(m^2+2)$ for experiment C (recall that the number of bidder variables for experiment B is 40 and that of C is m^2). If the variable is not selected to mutate, its value does not change. Otherwise, the new variable Var_i^{Mut} is computed by adding or subtracting a small value $a_i \cdot r_i$. Addition or subtraction is randomly chosen with probability *0.5*. r_i is set to be *0.1*(1-0)* because the range of bidders' chromosomes and sellers' chromosomes are all between 0 and 1, and the mutation range is fixed to be 10% of this range [Mühlenbein and Schlierkamp-Voosen 93][2]. a_i is computed with the equation $a_i = \sum_k u 2^{-k}$, where u is initially 0 and is randomly chosen to be 1 with probability *1/k* and the suggested k value is 16. On average, there will only be one u with value 1, which makes $a_i = 2^{-i}$ and this will

---

[1] The mutation rate is set to be 1/number of variables. In our experiments, bidders have a lookup table with 20 entries and each entry is treated as an independent variable. Sellers have 2 chromosomes, α and *P*; therefore, there are 22 variables in total and the mutation rate is set to be 1/22.

[2] Bidders' chromosomes contain variables representing bids between 0 and 1; sellers' chromosomes include probability of his action, α, which is between 0 and1, and the uniform-price offer, *P*, which is also between 0 and 1.

be the addition or subtraction to the original chromosomes [Mühlenbein and Schlierkamp-Voosen 93].

5.3 Detailed Steps

After the population is initialized, each bidder submits a bid according to its bidding function composite of a lookup table. Notice that, if the bidder's assumptions about the seller's values of α and P are correct, and the bidder can determine a best response to these values, this is indeed an optimal bid. The seller's fitness score is the total revenue over both stages, while that of the bidder is the individual payoff after both stages.

The bidder's payoff is calculated as the difference between its purchasing price — either its submitted bid at stage I or a fixed price offer — and its valuation. If the seller offers a fixed price higher than the bidder's valuation, the offer is rejected, and does not contribute to the seller's revenue, and the bidder's payoff is zero. Because the fitness function takes inputs from the chromosomes of both sellers and bidders, sellers and bidders can be viewed as two species that affect each other while evolving over generations. Figure 5-1 illustrates the detailed steps in one generation for experiment set A. Seven different experiments are conducted. For experiment B and C, the detailed steps are the same except for the chromosome encoding and the bidder valuation distribution.

As illustrated in figure 5-1, the number of generation differs for different experiments. For experiment set A, $y = 10{,}000$; for experiment set B, $y = 4{,}000$; for experiment C, $y = 5{,}000$.

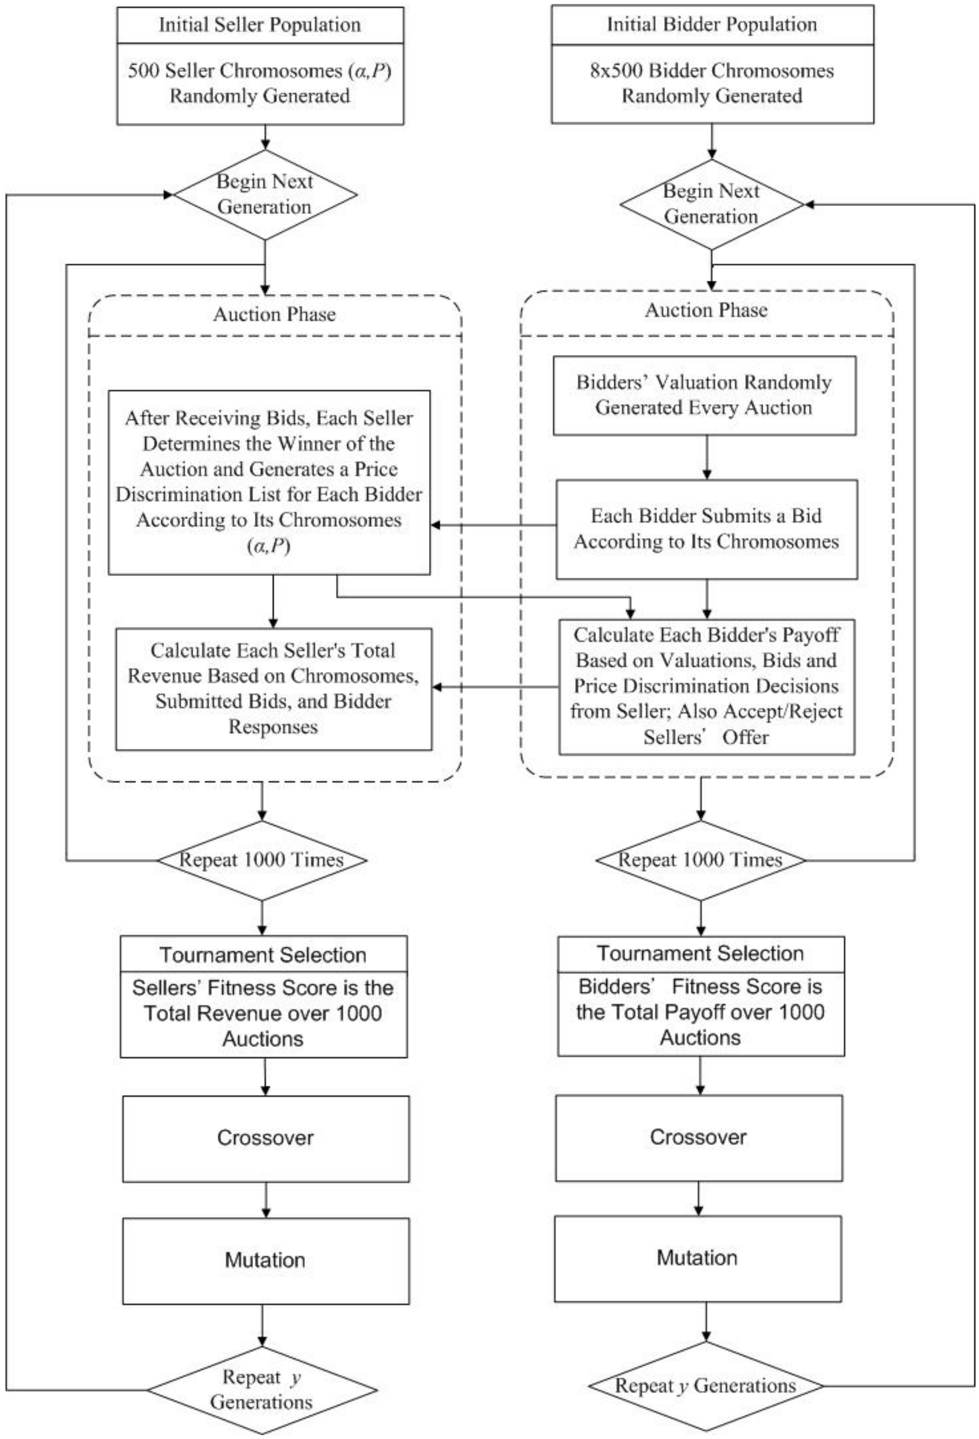

Figure 5-1: Detailed Process of Evolutionary Programming Experiments

5.4 Experiment Set A: Preliminary Results

This section describes preliminary experiments to determine some parameters that are used in later, more extensive experiments. We conduct an experiment of the independent first price auction, for which the Nash equilibrium strategy is well-known, and can be theoretically derived (see section 4.1). We explore–with 2, 4 and 8 bidders–whether evolutionary programming simulations can result in the correct equilibrium strategy. We choose 500 sellers, and 100 auctions per generation. Figure 5.2 compares the theoretical equilibrium strategy and the bidder chromosome obtained from experiments for 2 bidders. It shows that the experiments do provide correct results, but that these results contain some "noise". Results for 4 bidders and 8 bidders also provide similar, ``noisy'' results. We expect that some of the noise can be eliminated with an increase in the number of auctions per generation (a larger number of auctions results in a better approximation of the payoff due to different valuations among other bidders). However, some of the noise is also due to the existence of mutation, because of which the current strategy is slightly perturbed from the optimal strategy.

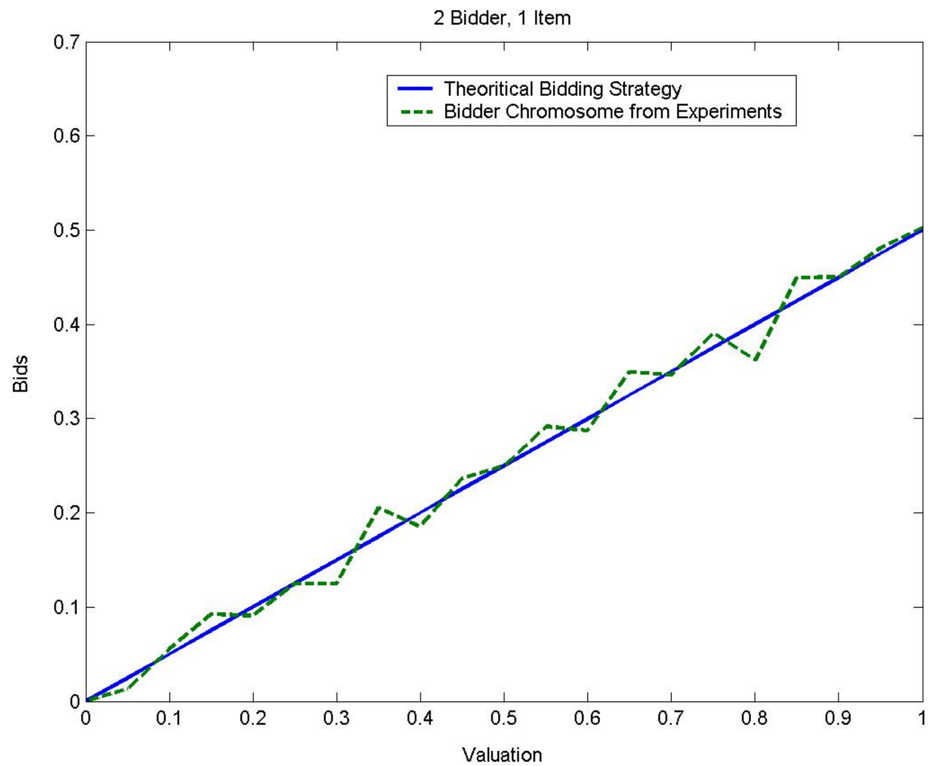

Figure 5-2: Experiment Results Compared to Theoretical Strategy

5.5 Experiment Set A: Simulation Results

5.5.1 Experiment One, Two and Three--n = 8 Bidders, *k = 2, 3,* and *4* Items

In all three experiments, the seller's α chromosome converges to $\alpha = 0.9999$ and the P chromosome does not converge. This is because when α is nearly 1, a second chance offer always occurs in Stage II; hence the fixed price, P, is never tested; and therefore does not converge. The bidding function (lookup table) for 2 items converges as in Figure 5.3.

We compare the analytically-obtained strategy derived in Chapter 4 with the experimental results. When $\alpha = 1$, the analytically-obtained strategy for valuation less or more than P is the same based on equation (1) in Chapter 4.

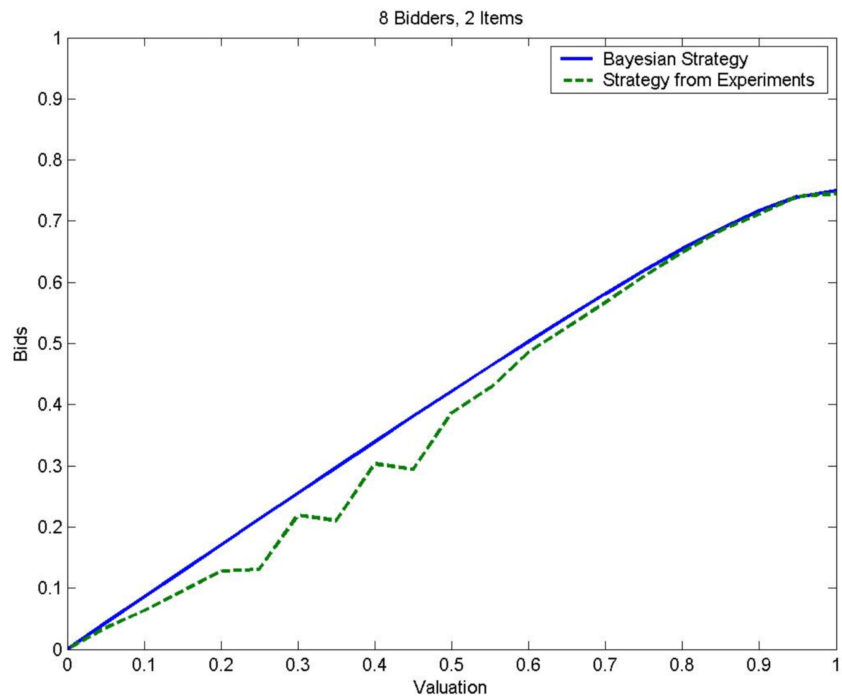

Figure 5-3: Strategy Comparison of Experiment One

Figures 5-3, 5-4 and 5-5 show both the analytically-obtained strategy and the bidding strategy obtained from experiment results for all three experiments.

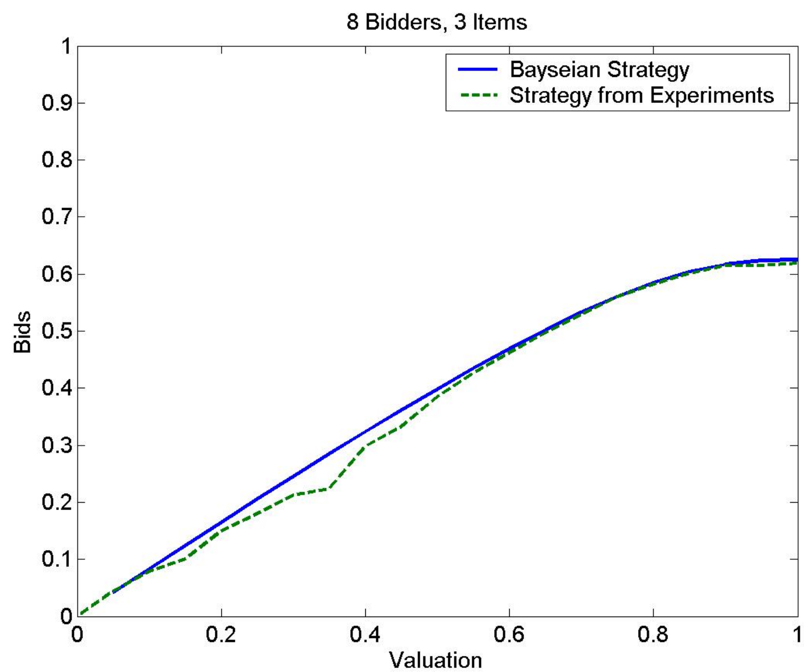

Figure 5-4: Strategy Comparison of Experiment Two

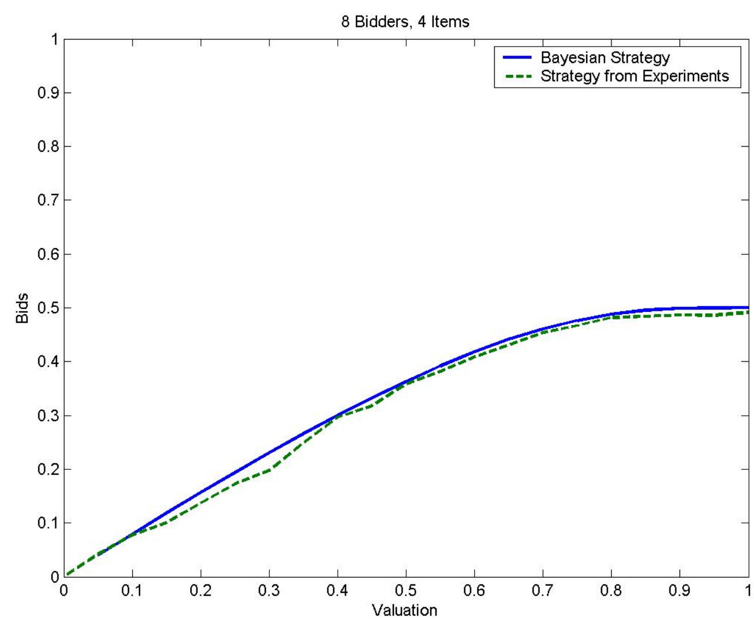

Figure 5-5: Strategy Comparison of Experiment Three

In all three experiments, it can be observed that even though the two strategies are not identical, the bidding strategy obtained using genetic simulations is similar to the analytically-obtained strategy; especially for high valuations. Solutions to the analytically-obtained strategies for these values of k are increasing, but not injective. A reasonable explanation is that because the bidders' fitness function is defined as *expected payoff*, and a non-zero contribution to the fitness function occurs only when the bidder wins, a higher percentage of the total payoff occurs for a high valuation, because a high valuation has a higher probability to win the auction with higher bids. When the bidders' fitness function is changed to $\frac{payoff}{H(x)+G(x)}$, figure 5-6 shows the preliminary results from an experiment that only conducts 200 auctions per generation after 500 generations. The bidder's expected payoff is denoted as $G(x)(x-b)+H(x)[\alpha(x-b)+(1-\alpha)(x-P)]$ from chapter 4. Because the seller's α converges to 0.9999, we can then rewrite the expected payoff to $G(x)(x-b)+H(x)(x-b)=[G(x)+H(x)](x-b)$. The modified fitness function is to optimize *x-b*. It is important to note that the modified fitness function cannot be applied to experiments where the optimal value of α is not known to be 1 because the expected payoff equation is different.

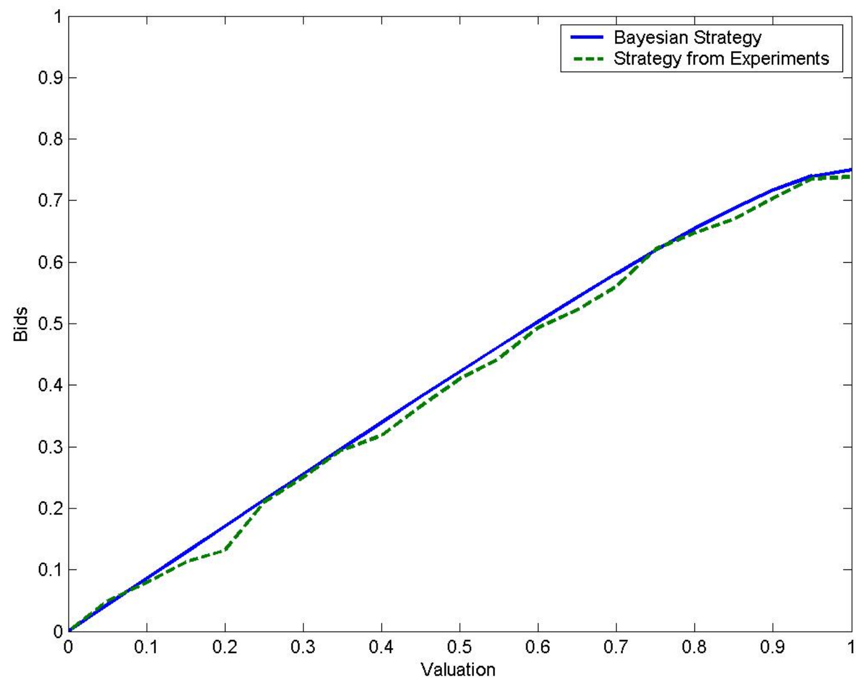

Figure 5-6: Bidding Function Comparison Using the Modified Fitness Function

We conduct a numerical analysis with the bidding function obtained from experiments with the modified fitness function. We first show that any different combination of P and α does not increase total revenue in figure 5-7, which provides average revenue over 10,000 auctions. In figures 5-8 and 5-9, we show that neither random over-bidding nor random under-bidding with any percentage range increases bidder payoff. Bidder No. 1 over-bids and under-bids while bidder Nos. 2 to 8 bid according to the bidding function. Therefore, we've show that near Nash equilibrium is obtained from the numerical analysis thus it is evolutionary stable. Deviation does not benefit. As noted in [Riechman 01], results obtained from genetic algorithms are not considered perfect Nash equilibrium. This is because, in a GA experiment, due to mutation, the entire population does not adopt the Nash symmetric strategy, but almost all of it does.

The main reason that we are exploring the difference between the original fitness function and the modified fitness function is because it shows that the modified fitness function provides similar results in fewer generations. This is crucial because we plan to conduct further experiments with a larger number of bidders. When the number of bidders increases, the number of auctions required to obtain a reasonable approximation of the average payoff increases exponentially[3] and the amount of time each generation takes to finish increases with it. With the modified fitness function, we can obtain analytically-obtained strategies from experiments with a smaller number of total generations, and a smaller number of auctions per generation.

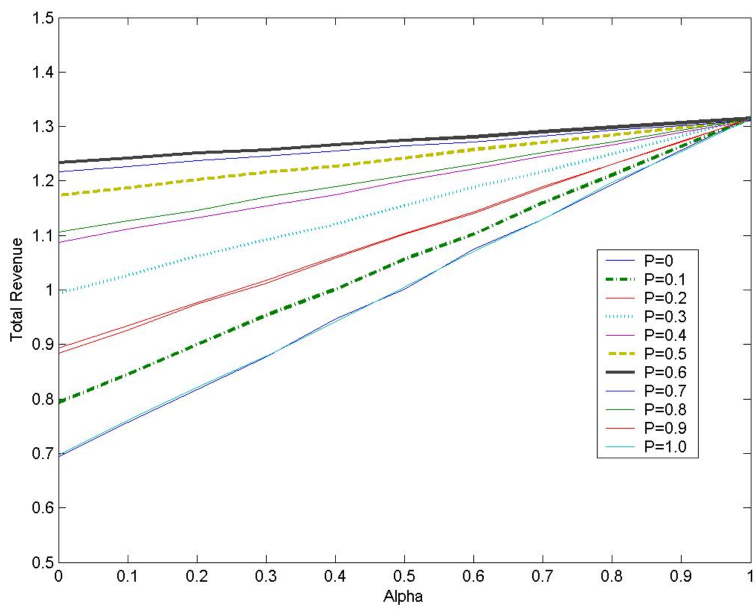

Figure 5-7: Total Revenue Comparison for Different *P* and α

[3] Because of $G(x) = x^{n-1}$, and *n* is the number of bidders.

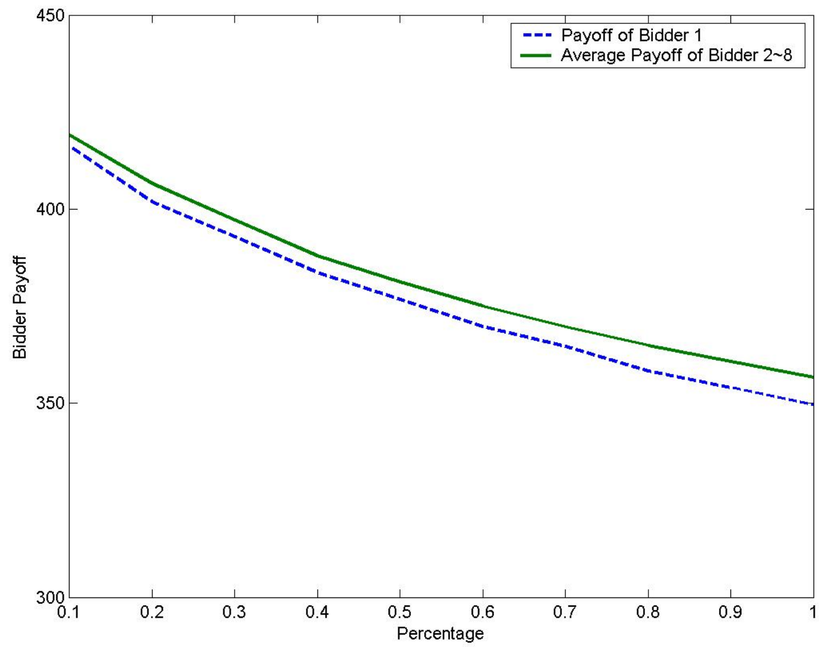

Figure 5-8: Bidder Payoff With Over-Bidding

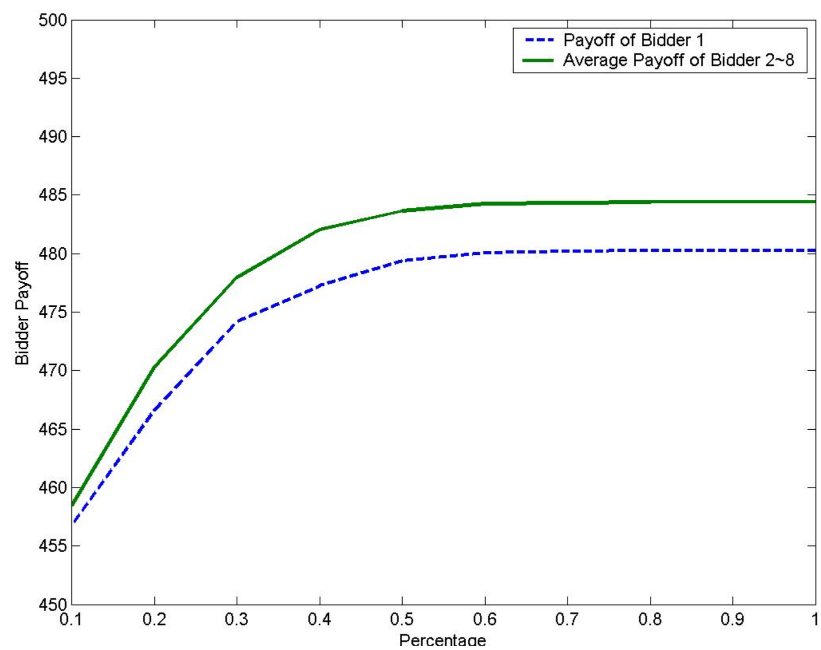

Figure 5-9: Bidder Payoff with Under-Bidding

5.5.2 *Experiment Four and Five--$n = 8$* Bidders, $k = 5$, and 6 Items

In the experiment with $k = 5$ items, the seller's value of α converges to 0.0015, and the value of P converges to 0.507. In the experiment with $k = 6$ items, the seller's value of α does not converge between 0 and 0.5, but the value of P converges to 0.4776. The results suggest that the seller's best response is randomizing between a second chance offer and a fixed price offer. The optimal fixed price offer is 0.507 and 0.4776 respectively for $k = 5$ and 6. Numerical analysis is conducted in a later section to show that this is consistent with a near Nash equilibrium strategy. Figure 5-10 shows the comparison between the bidding function obtained from the experiment with $k = 5$ items and analytically-obtained strategies with $\alpha = 0.0015$. Figure 5-11 shows the comparison between the bidding function obtained from the experiment with $k = 6$ items and analytically-obtained results with α equal to 0.1, 0.3 and 0.5 because α does not converge to a single value between 0 and 0.5. As shown in the figure, several analytically-obtained strategies are not monotonic increasing; hence the experimental outcomes are not expected to match the analytically-obtained strategies. The analytically-obtained strategies are not injective in both experiments.

In Section 5.7 we study the variation in the bidding strategy, across the population in a single experiment, as well as across experiments. We find that there is a very small variation across the population as well as across strategies for $k = 5$; this leads us to conclude that the genetic algorithm has converged to a symmetric equilibrium bidding strategy. On the other hand, for $k = 6$, we find a small variation (among bidding strategies averaged over the population) across experiments, but a larger variation across the population for a single experiment. Further, the variation across the population does not change significantly with an increase in the number of generations, leading us to conclude that the experiment does not converge to an equilibrium strategy; further, that the average strategy obtained is consistent across experiments. *We hence conclude that the strategy of Figure 5-7 is not the equilibrium bidding strategy for k = 6*, and that bidding strategies are likely to be mixed strategies, and hence that we need a different approach to parameterize the strategies. Details on the study of the variation may be found in section 5.7, and details on the study of mixed strategies may be found in section 5.9.

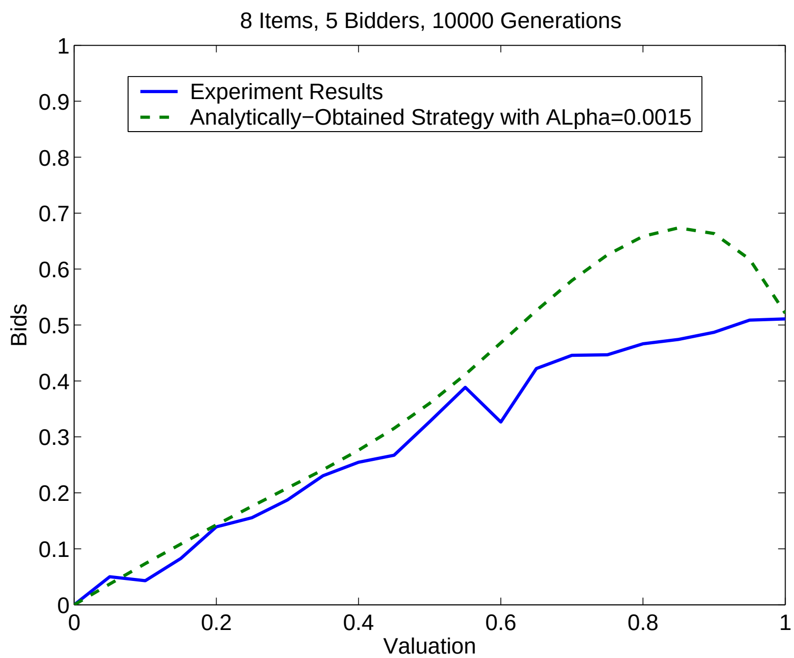

Figure 5-10: Strategy Comparison of Experiment Four

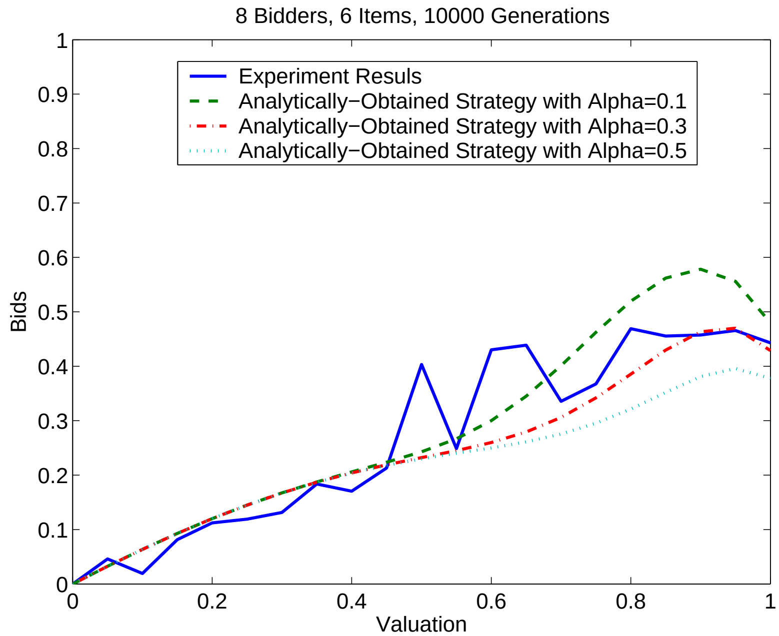

Figure 5-11: Strategy Comparison for Experiment Five

5.5.3 *Experiment Six and Seven—n = 8* Bidders, *k = 7* and *8* Items

In the experiment with k = 7 items, the seller's α chromosome converges to 0.09, and the P chromosome converges to 0.4684. Results suggested that seller's best response is to always make a fixed price offer at 0.4684. In the experiment with k = *8* items, the seller's α chromosome converges to 4.9e-07, and the P chromosome to 0.4809. The following figures show the comparison between bidding strategies obtained from the experiment and analytically-obtained strategies. In the experiment with k = *7* items, the difference between the experimental results and the analytically-obtained strategies appear to be because, in the GA simulation, because fixed-price offers are made to randomly chosen bidders, bidders with valuation greater than P can risk low bids when α is small. In the experiment with k = *8* items, the analytically-obtained strategy is not monotonic increasing.

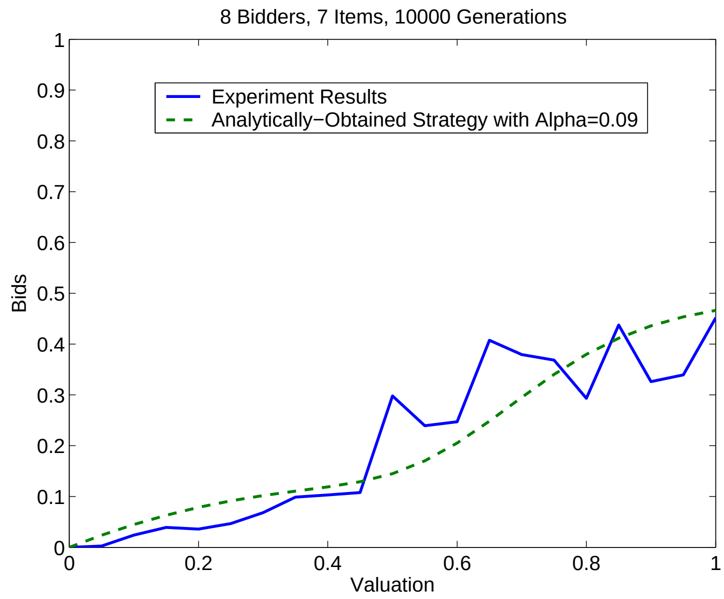

Figure 5-12: Bidding Strategy from Experiment Six

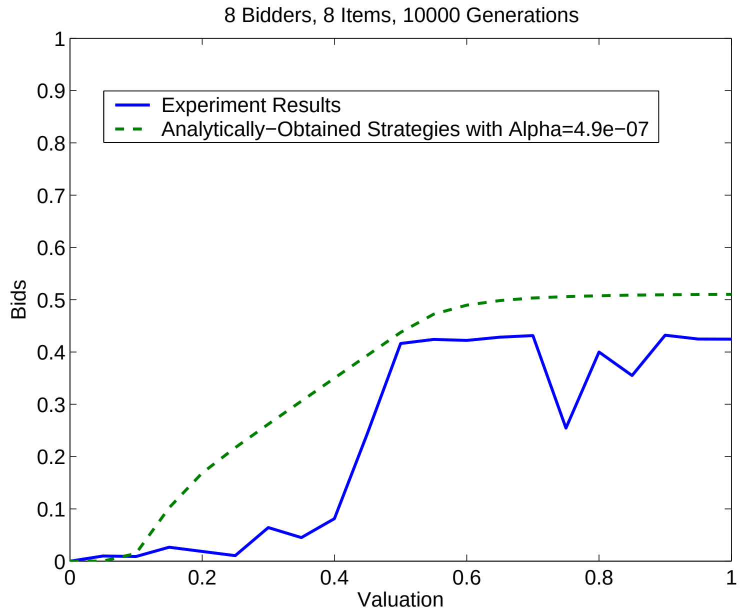

Figure 5-13: Bidding Strategy from Experiment Seven

We study the variance across the population and across experiments in section 5.7, and conclude, as with $k = 6$, that the bidding strategies in Figures 5-12 and 5-13 are not equilibrium bidding strategies, and that equilibrium bidding strategies for $k = 7$ and $k = 8$ are likely to be mixed strategies. We study mixed strategies in section 5.9.

5.6 Numerical Analysis

5.6.1 Demonstrate Near Nash Equilibrium for Experiments Four and Five

Bayesian strategies for experiment four have negative values and are not invertible as shown in Figure 5-10. To show that the experimental results reach Nash equilibrium, we conduct a numerical analysis from both the seller's side and the bidder's side. In experiment four, experimental results show that the seller's value of α does not converge, and that the seller's value of P converges to 0.507. In experiment five, it shows that the seller's value of α does not converge between 0 and 0.1, but that the seller's value of P converges to 0.4776.

5.6.1.1 Sellers' Side

For experiment four, we first show that any different combination of P and α does not increase total revenue in Figure 5-14. It also shows that with P fixed at 0.5, total revenue does not change for any α. This is average over 5,000 auctions. We show the same result for experiment five with α ranging from 0 to 0.1 in figure 5-15.

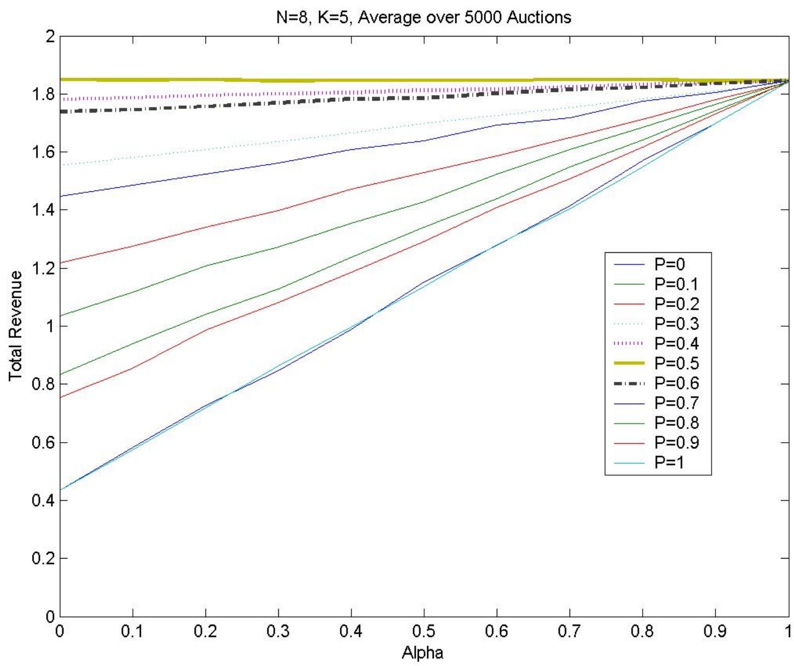

Figure 5-14: Total Revenue Comparison for Different α and P, Experiment 4

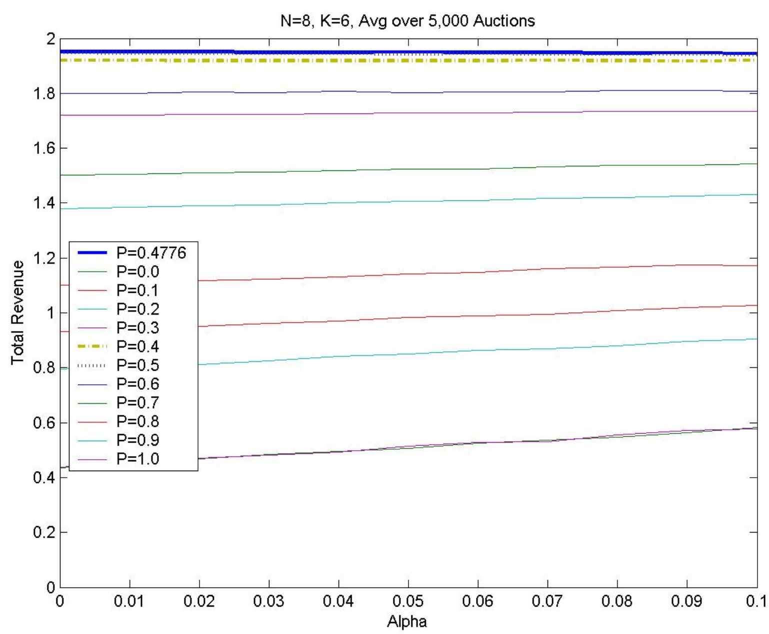

Figure 5-15: Total Revenue Comparison for Different α and P, Experiment 5

5.6.1.2 Bidders' Side

If the seller's α and P are the same as obtained from experiments, we show that over-bidding does not profit bidders in figure 5-16 and 5-17 for experiment four and five correspondingly. All the following numerical results are obtained from averages over 5,000 auctions as well. In figure 5-16 and 5-17, bidder No. 1 always over-bid a random percentage between 0 to 100%, but not exceeding its valuation, and bidder No. 2 to No. 8 bid according to the bidding function obtained from experiment results.

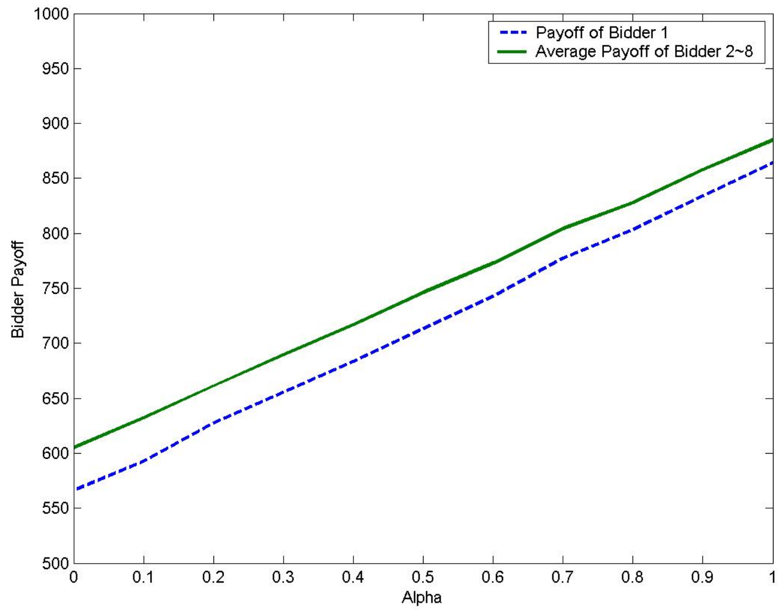

Figure 5-16: Bidder Payoff Comparison for Over-bid, Experiment 4

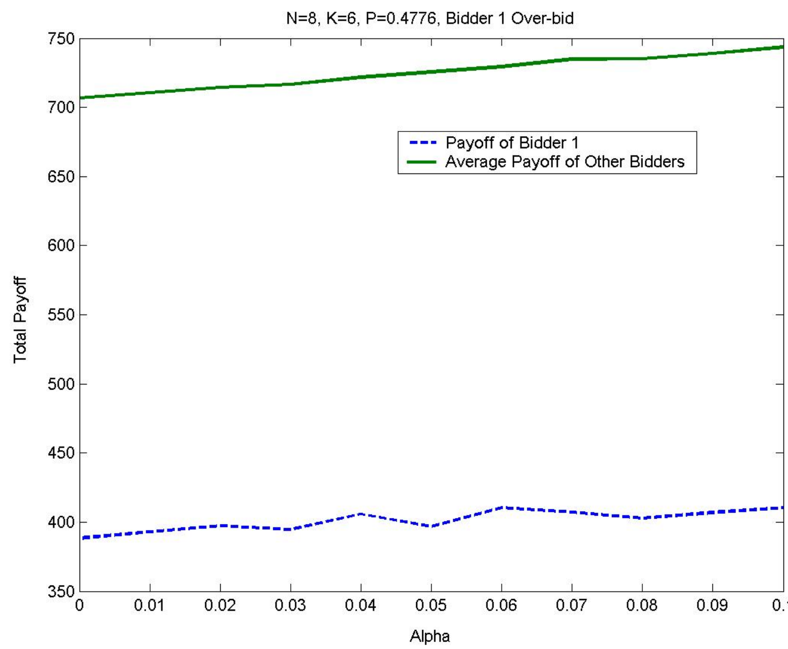

Figure 5-17: Bidder Payoff Comparison for Over-bid, Experiment 5

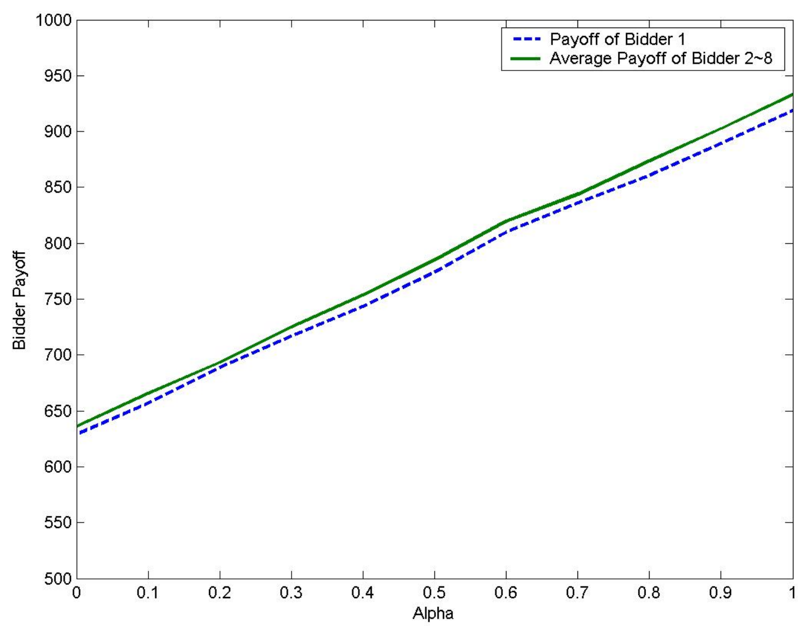

Figure 5-18: Bidder Payoff Comparison for Under-Bid, Experiment 4

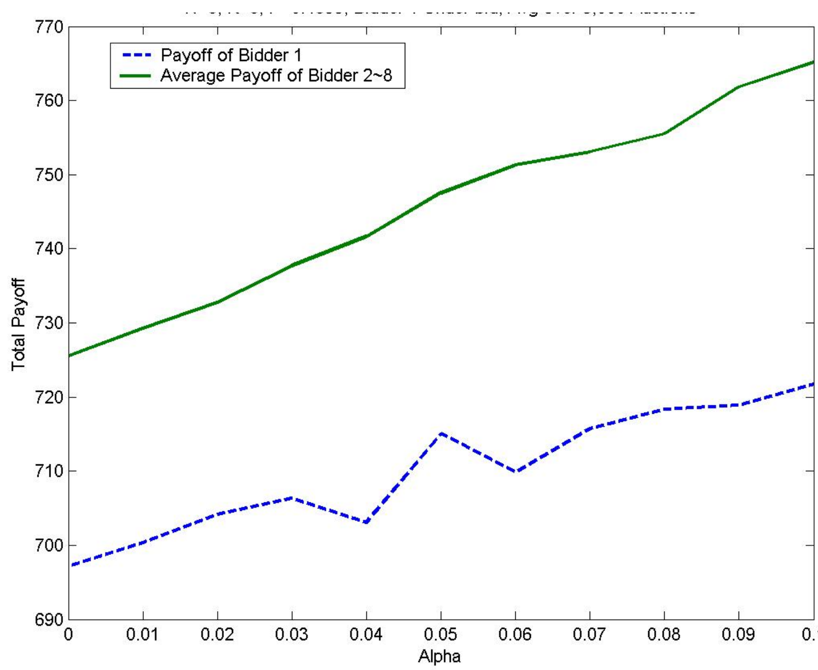

Figure 5-19: Bidder Payoff Comparison for Under-Bid, Experiment 5

We also show that under-bidding does not profit bidders in figures 5-18 and 5-19 for experiment four and five. In each auction, bidder No. 1 always under-bid a random percentage between 0 and 1, but no lower than zero, and bidder No. 2 to No. 8 bid according to the bidding function obtained from experiment results.

Furthermore, we show that over-bidding does not profit bidders even for a small percentage in figures 5-20 and 5-21. In each auction, bidder No. 1 always randomly over-bid in a 10% range and bidder No. 2 to No. 8 bid according to the bidding function obtained from experiments.

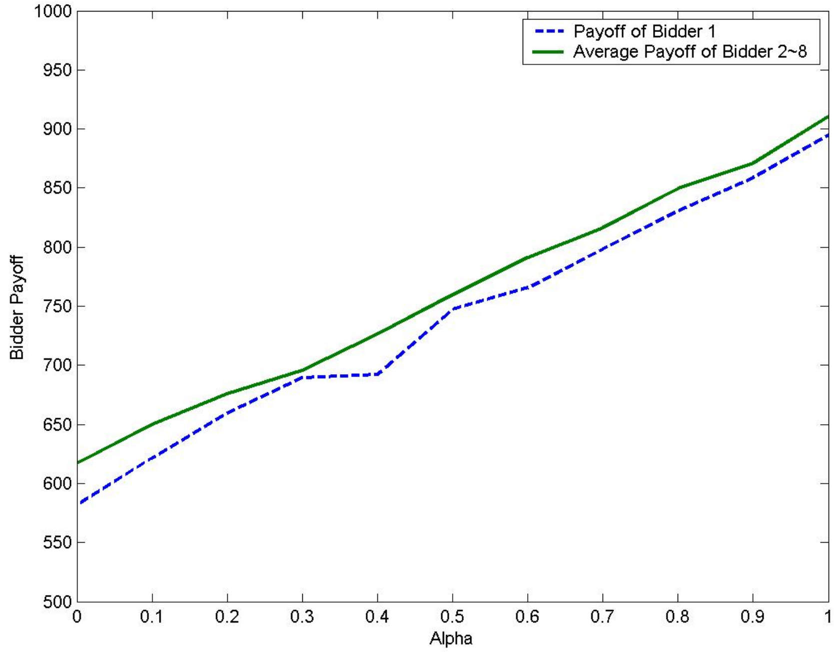

Figure 5-20: Bidder Payoff Comparison for Over-Bid in a Small Range, Experiment 4

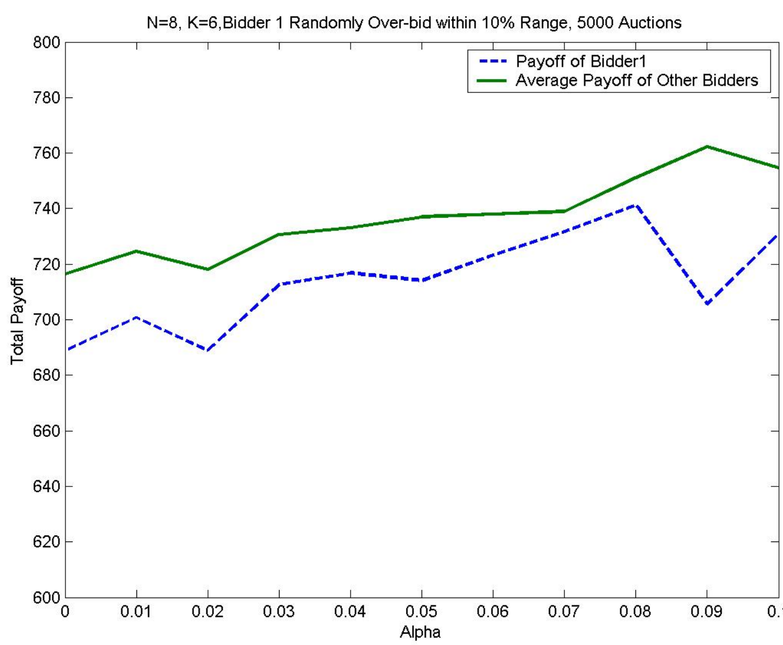

Figure 5-21: Bidder Payoff Comparison for Over-Bid in a Small Range, Experiment 5

Similarly, in figures 5-22 and 5-23, we show that under-bidding does not profit bidders even for a small percentage. In each auction, bidder No. 1 always randomly under-bid in a 10% range and bidder No. 2 to No. 8 bid according to the bidding function obtained from experiments. We show that neither over-bidding nor under-bidding in a small range results in higher bidder payoff even in a small range; therefore, deviation from the bidding function obtained from experimental results does not benefit bidders. We also show that deviation from α and P for the seller does not benefit sellers. We can now conclude that results in experiment four and five reached Nash equilibrium because it consists of mutual best responses for sellers and bidders.

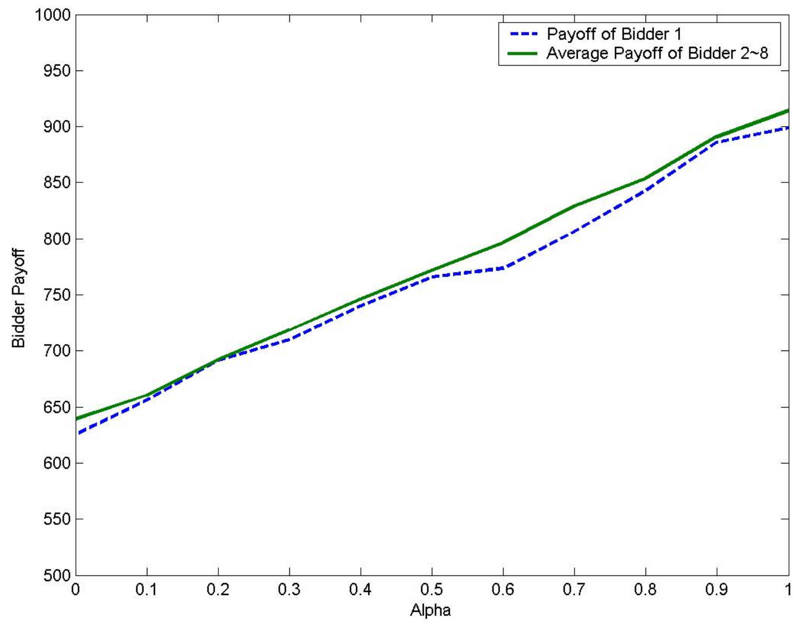

Figure 5-22: Bidder Payoff Comparison for Under-Bid in a Small Range, Experiment 4

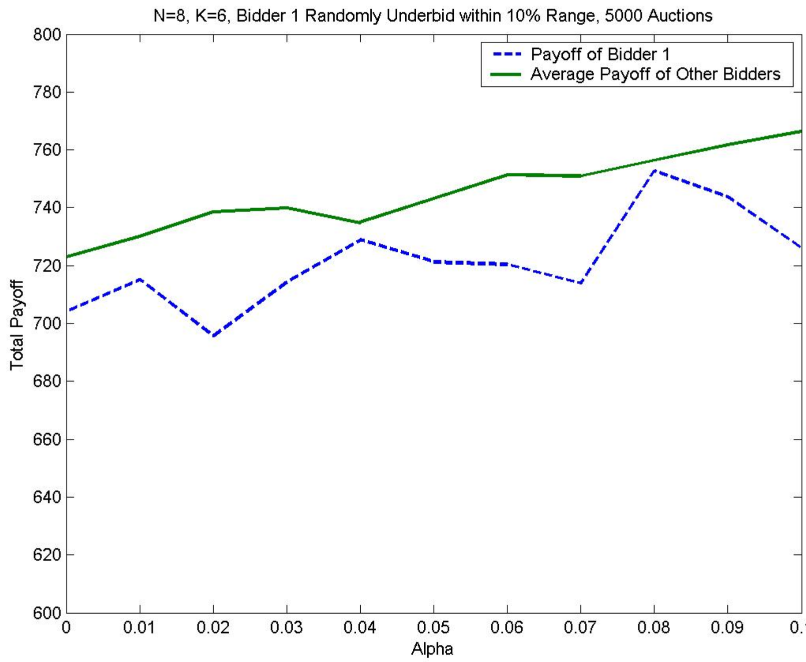

Figure 5-23: Bidder Payoff Comparison for Under-Bid in a Small Range, Experiment 5

5.6.2 Show Nash Equilibrium for Experiments Six and Seven

Similar to experiments four and five, Bayesian strategies for experiment six and seven also have negative values and are not monotonic increasing. To show that the experimental results reach Nash equilibrium, we conduct a numerical analysis from both sellers' side and bidders' side. In experiment six, experimental results show that the seller's value of α converges to 0.001, and the seller's value of P converges to 0.4684. In experiment seven, we see that the seller's value of α converges to 0.001, and the seller's value of P converges to 0.4809. We obtain the same results for experiment seven; in the following section we only show the figures from experiment six.

5.6.2.1 Sellers' Side

We first show that any different combination of P and α does not increase total revenue in figure 5-24. It also shows that with P fixed at 0.4684, the total revenue reaches its highest at α near 0. The values are averaged over 5,000 auctions.

5.6.2.2 Bidder's Side

If the seller's values of α and P are the same as obtained from experiments, we show that over-bidding does not profit bidders in figure 5-25. All the following numerical results are obtained by averaging over 5,000 auctions. In each auction, bidder No. 1 always over-bid a random percentage and bidder No. 2 to No. 8 bid according to the bidding function obtained from experiment results. We show that for any percentage, over-bidding does not result in higher payoff.

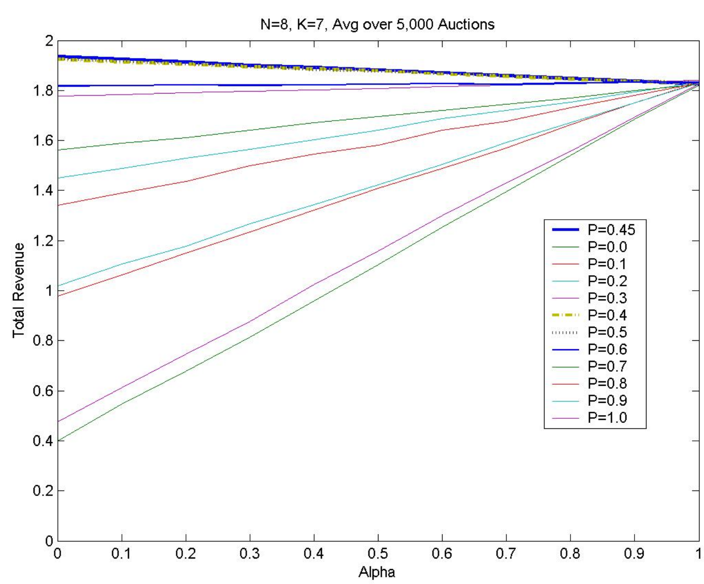

Figure 5-24: Total Revenue Comparison for Different α and P

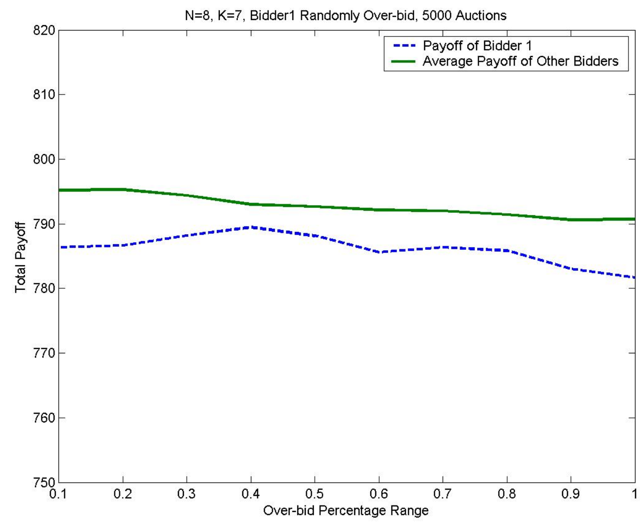

Figure 5-25: Bidder Payoff Comparison for Over-Bid

Similarly, in figure 5-26, we show that under-bidding does not profit bidders for any percentage. In each auction, bidder No. 1 always under-bid and bidder No. 2 to No. 8 bid according to the bidding function obtained from experiment results

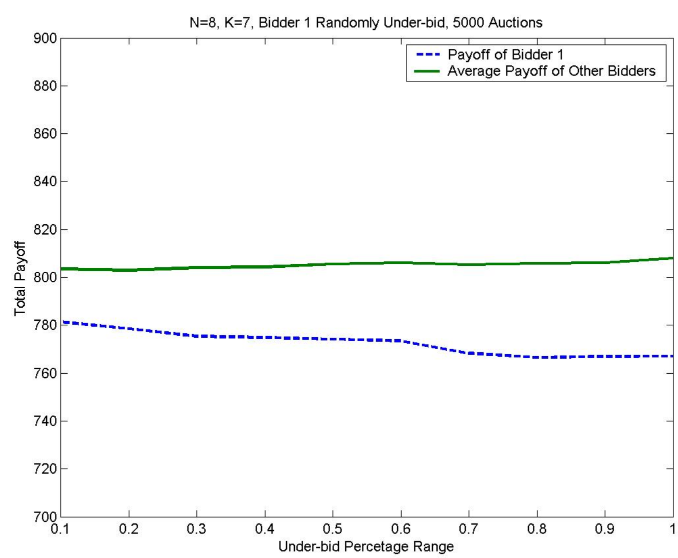

Figure 5-26: Bidder Payoff Comparison for Under-Bid

Based on the numerical analysis above, neither deviation for bidders nor deviation for sellers benefits them. It is shown that Nash equilibrium is reached because it is a mutual best response for both sellers and bidders.

5.7 Mixed Strategy Analysis

As pointed out in [Riechmann 01], the variance of the chromosome across the population in genetic algorithm simulations can be used as a measure of convergence. The smaller the variance is, the more converged the genetic population is. That is under the assumption that the solution is a pure, non-probabilistic strategy. If there is no solution that is a pure non-probabilistic strategy, the the population variance will remain high.

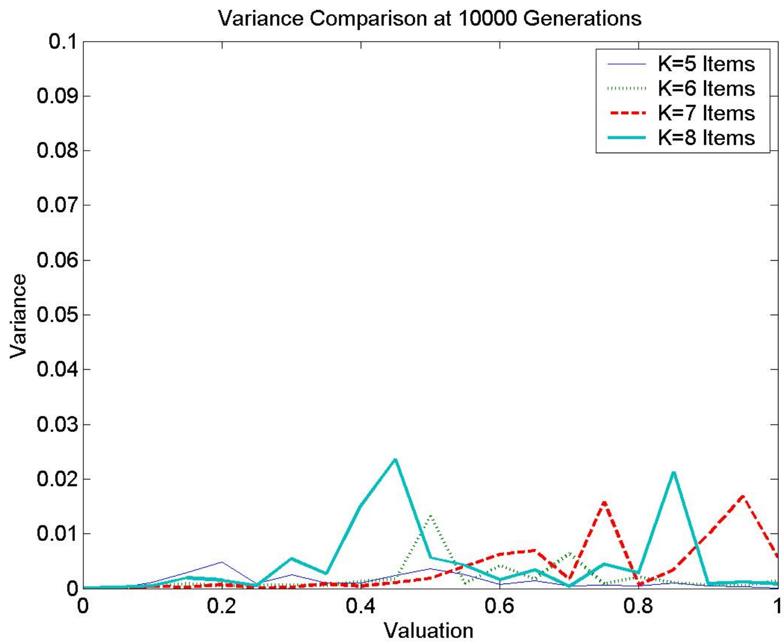

Figure 5-27: Comparison of Variance of Round 1

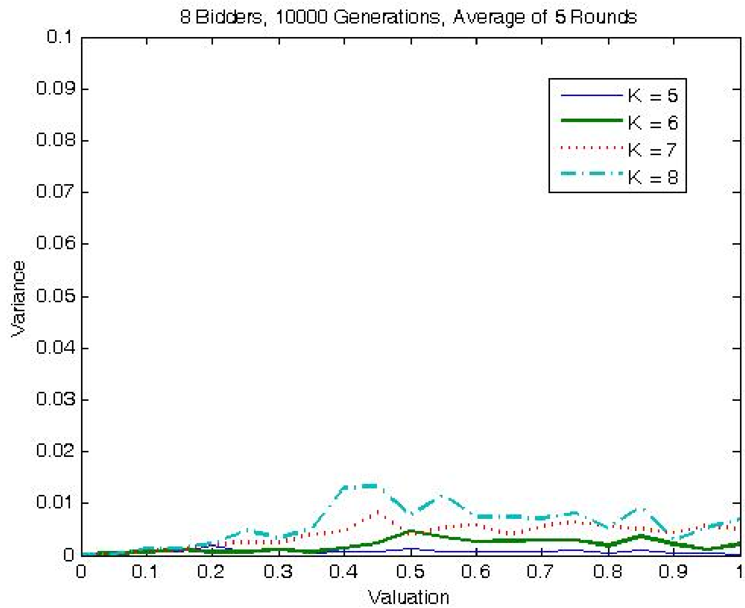

Figure 5-28: Comparison of Average Variance over 5 Rounds

We studied the variance across the bidder population, as well as across experiments. We notice that the population variance for the bidder chromosome is considerably larger for $k = 6$, 7 and 8 than that for $k = 5$, see Figure 5-27 for the population variance of the first round and Figure 5-28 for average population variance of 5 rounds of experiments. We also observe the variance over 6,000, 7,000, 8,000, 9,000 and 10,000 generations for $k = 6$, 7 and 8, and determine that it does not change significantly. We hence conclude that the population has achieved evolutionary stability; however, the larger variance implies that it is possible that the strategy is not a deterministic strategy, but is, instead, a mixed strategy, which our chromosome encoding scheme does not capture accurately.

Consider the bidding strategy result of a single experiment (that is, the average value of the bidder chromosome across the population) as a vector of random variables (the value of the bid for each valuation is a random variable). We carried out five such experiments, obtained five vectors of instances of the random variables, and used these samples of the random variables to estimate their variances. The estimated variance, as a function of valuation, is plotted in figure 5-29 for *k* = *5*, *6*, *7*, and *8*. As plotted in figure 5-29, we observe that each of these values is very small for *k* = *5*, *6*, *7* and *8*. We hence conclude that the observed results are indeed repeatable, and hence that our conclusions, that we have obtained bidding strategies for *k* = *5* but not for *k* = *6*, *7* and *8* are correct. In experiment B, we change the encoding scheme of the bidder chromosomes to allow mixed strategies. It is important to note that the encoding scheme still allows pure strategies. Details are discussed in section 5.9.

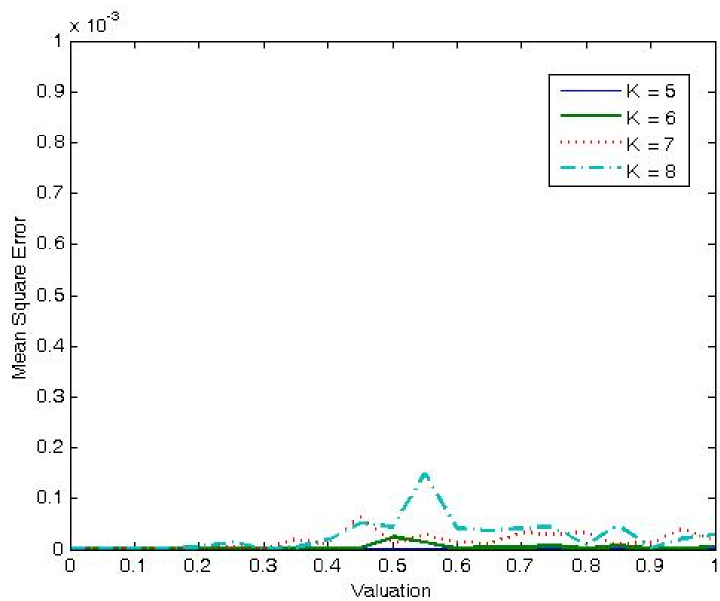

Figure 5-29: Unbiased Estimate of Variance in Bidding Strategy across Experiments

5.8 Experiment Set A: Summary

We summarize our experimental results as follows, in this 8-bidder-2-stage game:

1. *k = 2, 3, and 4*: The seller's best response is to price discriminate. The bidder's best response is consistent with the Bayesian strategy. This means that the seller prefers the privacy-infringing choice and the best the bidders can do is to bid the Bayesian strategy. Rational behavior does not decrease bids enough to deter seller from using price discrimination.
2. *k = 6*: The seller's optimal strategy is to randomize its action between price discrimination and a uniform-price offer. Bidders have a bidding strategy that is different from the Bayesian one and is likely to be a mixed strategy. This implies that the seller is indifferent about privacy-protection and privacy-infringing choices.
3. *k = 5, 7,8*: The seller's best response is to always make uniform-price offers and the bidder's best response is pure strategy for *k = 5*, and is likely to be a mixed, randomized strategy for *k = 7, 8*. This means that sellers prefer the privacy-protecting choice, and, for *k = 7, 8*, part of the bidder's best response is between dropping out by bidding zero, or revealing valuation by bidding its valuation.

To summarize experiment set A, when first-price auction is adopted at stage I, it is clear that different degrees of item scarcity make a difference for the seller's choice between privacy protection mechanism and privacy infringement mechanism. Bidders respond by changing their bidding strategy for different degrees of privacy invasion.

5.9 Experiment Set B: Mixed Strategy Experiment

From experiment set A, we observe that it is highly possibly to have a mixed strategy when the bidders' chromosome has high variance over the population. In this experimental set, we change the design of bidder chromosomes to analyze possible mixed strategies. In the previous two sets of experiments, the bidder's chromosome is encoded as a valuation-to-bid lookup table. Each valuation has a bid. It is harder to reverse engineer the composite of possible mixed strategies with only a valuation-to-bid mapping, because a mixed strategy is composed of probabilities *and* bids. For example, if a mixed strategy contains two strategies with 50% probability each: bid 0 and bid 1. It means that half of the time a bid 0 is submitted and a bid 1 otherwise. To better represent a mixed strategy, we encode the bidder chromosome as an *m-by-m* matrix where *m* is the number of bids that can be submitted. Each element within this matrix represents the probability of this bid. For example, the bidder chromosome can be illustrated in the following table:

Table 5-1: Bidder Chromosome Example

	Bid 0	Bid 0.5	Bid 1.0
Valuation 0	0%	0%	0%
Valuation 0.5	50%	50%	0%
Valuation 1.0	30%	60%	10%

In this example, with valuation *1.0, 30%* of the time the bidder submits a *0* bid and *60%* of the time a *0.5* bid is submitted. In this set of experiment, bidders are not allowed to submit a bid that is higher than its valuation; therefore the element *(2, 3) = 0%* because bidder with valuation *0.5* cannot submit a *1.0* bid. The idea of encoding bidder chromosome with a *m-by-m* matrix is to convert this problem from one where the bid is drawn from the continuous solution space *[0, 1]* to one where it is drawn from a discrete solution space of a few possible bids. In this set of experiments, the first-price auction and the uniform distribution are adopted. Bidder valuation ranges from *0* to *1* with equal probability.

5.10 Experiment Set B: Simulation Results

5.10.1 Experiment One--, $n = 8$ Bidders, $k = 6$ Items

We conduct two versions of this experiment. In the first version, the bidder chromosome is encoded by a *5-by-5* matrix. It means that each bidder can have 5 types of valuation: *0, 0.25, 0.5, 0.75,* and *1.0*. Each bidder is allowed to submit 5 different bids with different probability. No bid higher than valuation is allowed to be submitted. We obtain the following bidder chromosome matrix from the experiment after 4000 generations with 1000 auctions in each generation:

Table 5-2: Bidder Chromosome from Experiment One

	Bid 0.0	Bid 0.25	Bid 0.5	Bid 0.75	Bid 1.0
Valuation 0.0	100%	0%	0%	0%	0%
Valuation 0.25	3.31%	96.69%	0%	0%	0%
Valuation 0.5	3.82%	86.89%	9.29%	0%	0%
Valuation 0.75	2.94%	6.77%	89.87%	0.43%	0%
Valuation 1.0	2.78%	4.95%	91.93%	0.23%	0.12%

In this matrix we cannot observe an obvious mixed strategy because each type of valuation has a mapping bid with very high probability. It is more likely to be a pure strategy instead.

We then conduct the second version of the experiment with an *11-by-11* matrix for the bidder chromosome. This experiment also has 4,000 generations and 1,000 auctions per generation. We plot the bidder chromosome in figure 5-30. For valuation lower than 0.5, there does not exist an obvious mixed strategy. We extract the bidder chromosome of valuation higher than 0.5 and plot it in figure 5-31. We can observe that there are two major bidding strategies exhibited in the plot: bidding 0.5 and a less than 0.5 bid with different probabilities.

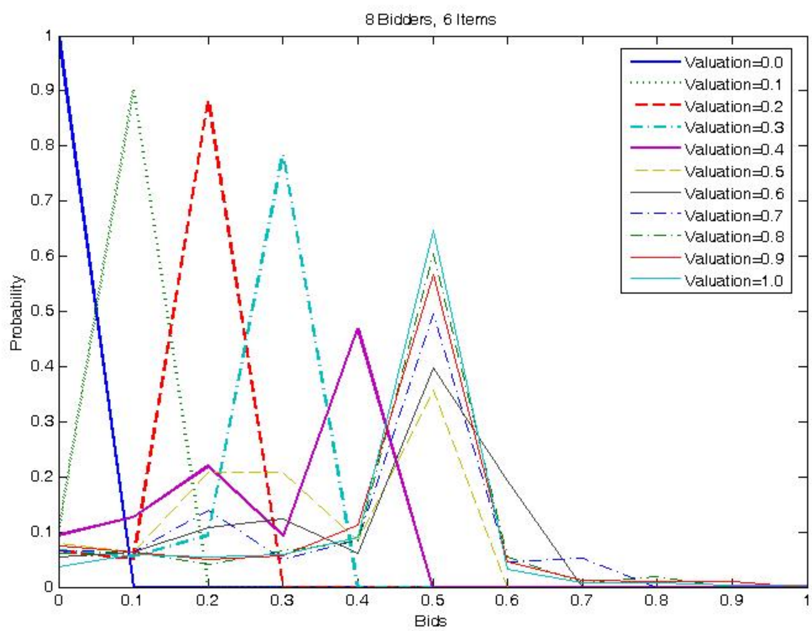

Figure 5-30: Bidder Chromosome from Experiment One

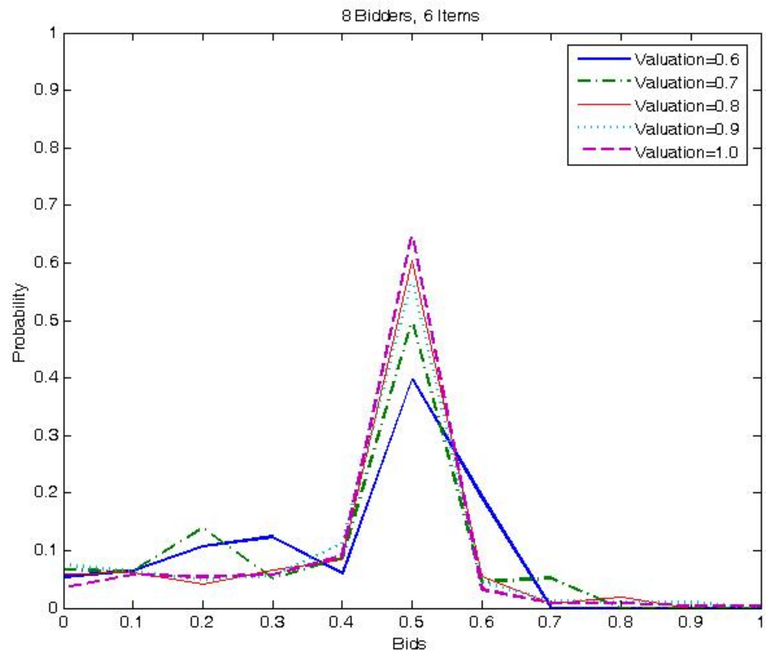

Figure 5-31: Bidder Chromosome from Experiment One, Higher Valuation

5.10.2 Experiment Two—n = 8 Bidders, *k = 7* Items

Similar to experiment one, we conduct two versions of the experiment with 8 bidders, $k = 7$ items, the uniform distribution, and the first-price auction. The following table illustrates the experimental results when the bidder chromosome is encoded as a 5-by-5 matrix.

The experimental results are different from experiment one, but, as in experiment one, no obvious mixed strategy is observed at this level of sampling of the bid/valuation space. We further conduct the second version of the experiment with an *11-by-11* matrix. The following two figures illustrate the bidder chromosome obtained from experiment.

Table 5-3: Bidder Chromosome from Experiment Two

	Bid 0.0	Bid 0.25	Bid 0.5	Bid 0.75	Bid 1.0
Valuation 0.0	100%	0%	0%	0%	0%
Valuation 0.25	83.71%	16.29%	0%	0%	0%
Valuation 0.5	1.7%	94.61%	3.7%	0%	0%
Valuation 0.75	1.13%	94.14%	3.98%	0.75%	0%
Valuation 1.0	1.48%	94.20%	3.80%	0.3%	0.22%

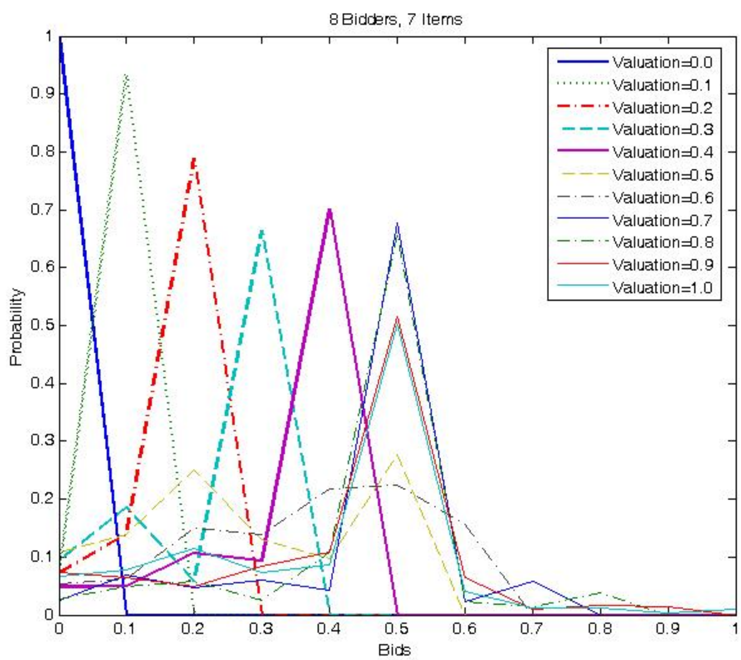

Figure 5-32: Bidder Chromosome from Experiment Two

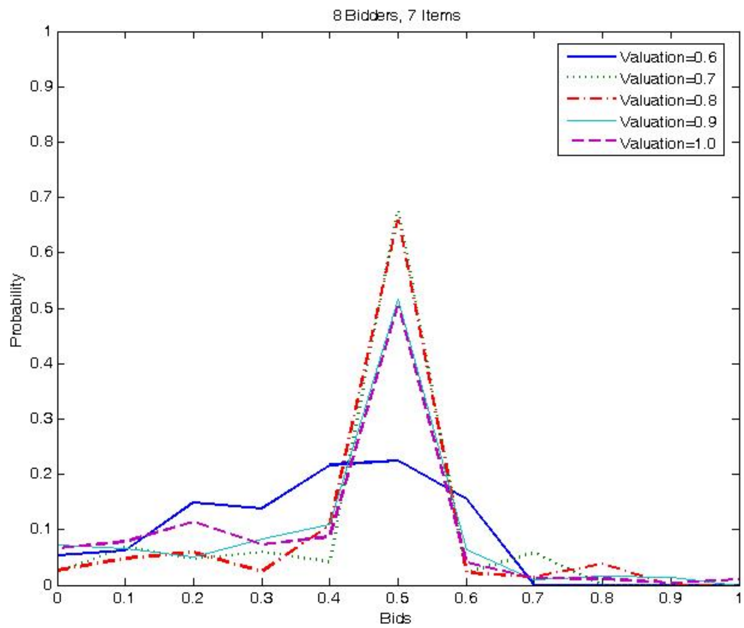

Figure 5-33: Bidder Chromosome from Experiment Two, Higher Valuation

5.10.3 Experiment Three—n = 8 Bidders, *k = 8* Items

Similar to experiments one and two, we conduct two versions of this experiment with $n = 8$ bidders, $k = 8$ items, the uniform distribution and the first-price auction. The following table illustrates the experimental results when the bidder chromosome is encoded as a 5-by-5 matrix.

The experimental results are different from those of experiment one, are similar to those of experiment two; no obvious mixed strategy is observed. On conducting the second version of the experiment with an *11-by-11* matrix, we obtain bidder chromosomes illustrated in Figures 5-34 and 5-35.

Table 5-4: Bidder Chromosome from Experiment Three

	Bid 0.0	Bid 0.25	Bid 0.5	Bid 0.75	Bid 1.0
Valuation 0.0	100%	0%	0%	0%	0%
Valuation 0.25	86.54%	13.46%	0%	0%	0%
Valuation 0.5	1.51%	94.76%	3.73%	0%	0%
Valuation 0.75	1.22%	95.24%	2.91%	0.63%	0%
Valuation 1.0	0.99%	96.41%	2.03%	0.35%	0.22%

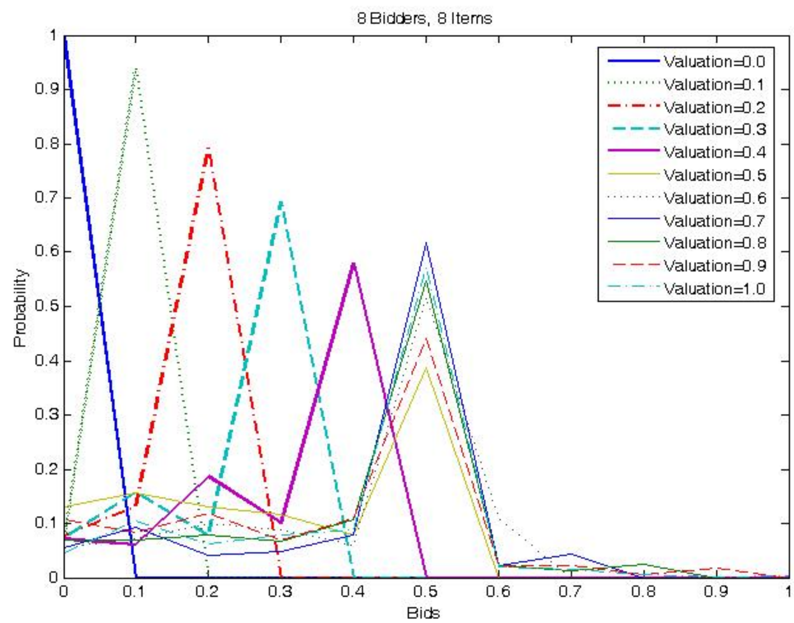

Figure 5-34: Bidder Chromosome from Experiment Three

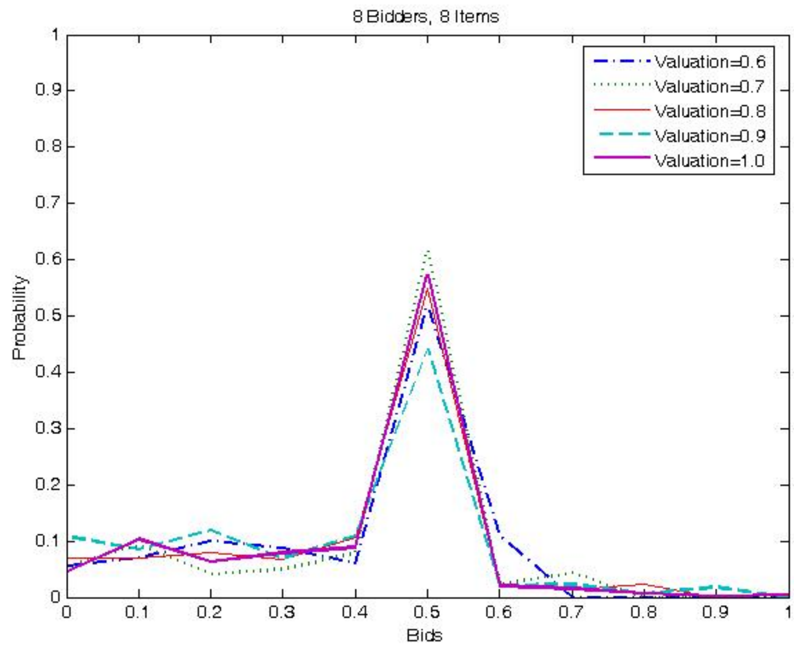

Figure 5-35: Bidder Chromosome from Experiment Three, Higher Valuation

In this experiment, from figure 5-35, we can observe it is likely to have two different mixed strategies: bidding 0.5 and bidding below 0.5 uniformly. The distribution of the bidding strategy below 0.5 is more likely to be uniform than in experiments one and two.

5.11 Experiment Set B: Summary

We summarize our experimental results as follows, when analyzing the mixed strategy:

1. Granularity matters: we have to conduct two versions of experiments with different granularity because the first version changes the nature of the game. We have to be careful about the observations we make with different granularity.
2. When converting the game into a discrete solution space, we cannot take the solution from the discrete version and conclude that the same solution also applies

to the continuous version. However, we can make better observations when the granularity is right. One way to find the right granularity is iteratively dividing the continuous space into more discrete actions; when the results don't change, we can know that this is the minimum granularity to mimic the game of continuous space. In our experiment, we have also conducted experiments with a *21-by-21* matrix as bidder chromosome. The result is not different from that of an *11-by-11* matrix as. We thus conclude that 11-by-11 matrix is the minimum granularity for this problem.

CHAPTER 6: EXPERIMENTAL DATA FROM *EBAY*

In this chapter, we describe the genetic algorithm simulation results of experiment set C. It takes input from real data on eBay as well as from a uniform distribution between the highest and lowest valuations estimated from eBay data. We also describe the implementation of a software bidding tool for bidders to adjust bidding strategies considering the possibility of a single second chance offer (that is, $k = 2$) after auctions end. It is developed based on our model and genetic algorithm experiments. In the following sections, the experimental data is presented, and the features, the platform and the language of the bidding tool are documented.

6.1 Experiment Set C: Bidder Valuation Distribution from eBay Data

For experiment set C, we collect real auction data from eBay to obtain the valuation distribution for bidders. We conduct experiments for both the real data distribution and the uniform distribution. There are certain constraints about gathering auction data from eBay. First, only data from auctions completed in the past two weeks are publicly available; second, a keyword search is needed to gather all auctions for the same item. Keyword "Wii 14" was used to pull all the auctions of a unique bundle of the Wii gaming console, two remote controllers and 14 games. We chose this item because it has a low inventory in retail stores and is thus very popular on eBay.

After gathering raw data, we clustered all the auction data by number of bidders. We observed that the number of auctions with 8 and 13 bidders was in the hundreds, large enough to make an accurate estimate of *F(x,)* the cumulative distribution function. For the uniform distribution, $F(x) = x$. The difficulty of computing *F(x)* from eBay data is the bids of an auction only show the valuation of all bidders except the winner. Because the highest valuation, x, is unknown, we have to compute an estimate of *F(x)* as follows:

- Step 1: for each auction, calculate $F_2(x)$:

$$F_2(x) = \frac{number\ of\ 2nd\ highest\ bidders\ below\ x}{total\ number\ of\ 2nd\ highest\ bidders(number\ of\ auctions)}$$

- Step 2: from $F_2(x)$ compute *F(x)* based on order statistics, where *n* is the number of bidders: $F_2(x) = nF(x)^{n-1}\left[1 - F(x)\right] + F(x)^n$

Results from previous two steps are pairs of *F(x)* and *x*; a polynomial function can be obtained with least square approximation using Matlab. We plot the result in the following figure.

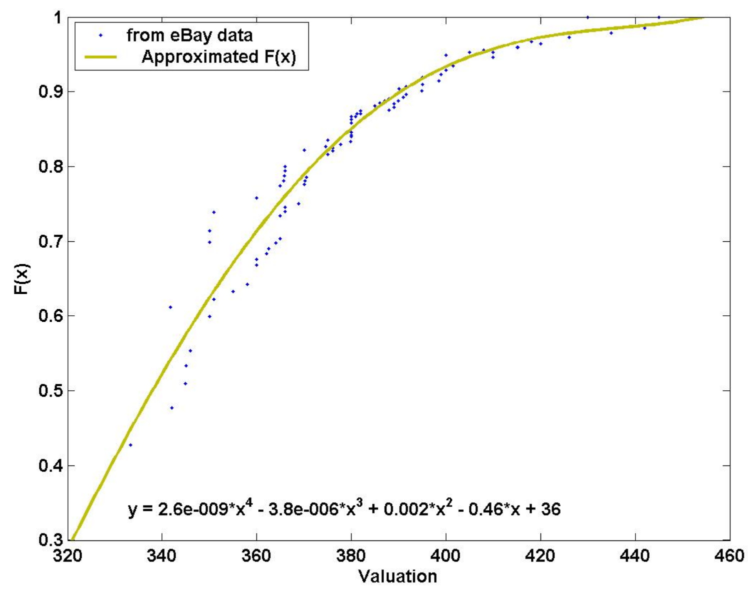

Figure 6-1: *F(x)* from Least Square Approximation

Seven experiments are conducted with n = 8 bidders each and number of items k =2, 3, 4, 5, 6, 7 and 8. Each experiment has two versions—version one has bidder valuation generated based on $F(x)$ obtained from real eBay auctions; version two is run with the bidder valuation set to be uniform (which would result in a linear function $F(x)$).

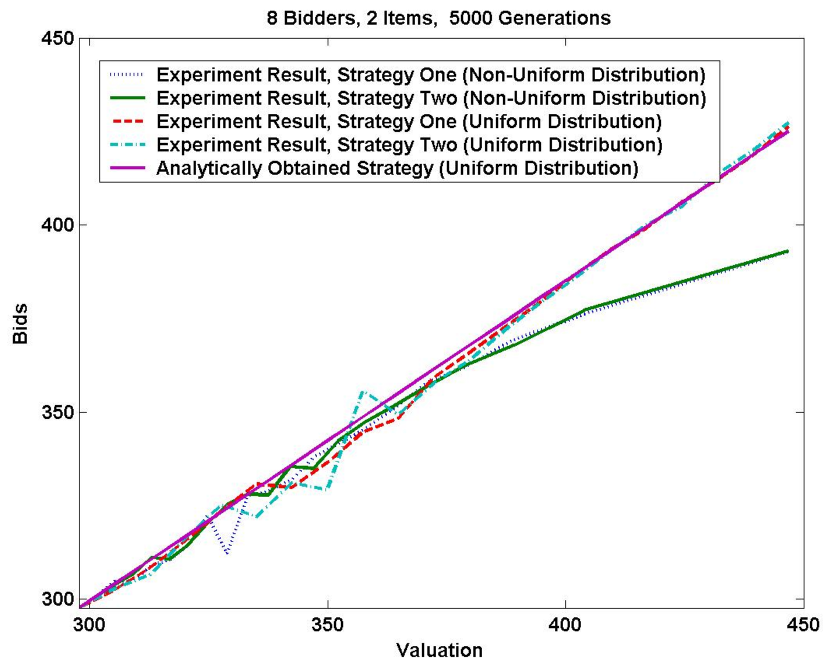

Figure 6-2: Bidding Strategy from Experiment One

6.2 Experiment Set C: Simulation Results

6.2.1 Experiment One and Two --n = *8* Bidders, k = *2*, and *3* Items

In all three experiments, the seller's α chromosome converges to α = *0.9999* and the P chromosome does not converge. This is because when α is nearly 1, a second chance offer always occurs in Stage II; hence the fixed price, P, is never tested; and therefore does not converge.

We compare the analytically-obtained strategy derived from Chapter 4 with the experimental results. When $\alpha = 1$, the Bayesian strategy for valuation less or more than P is the same based on equation (1) in Chapter 4. Figures 6-2 and 6-3 show both the analytically-obtained strategy and the bidding strategy obtained from experimental results for both experiments. For both the uniform and non-uniform distributions, the two lookup tables are almost identical. This implies that there is no strategy, in either experiment, that is a mix of two strategies, and, additionally, that a pure strategy is a symmetric equilibrium for both distributions.

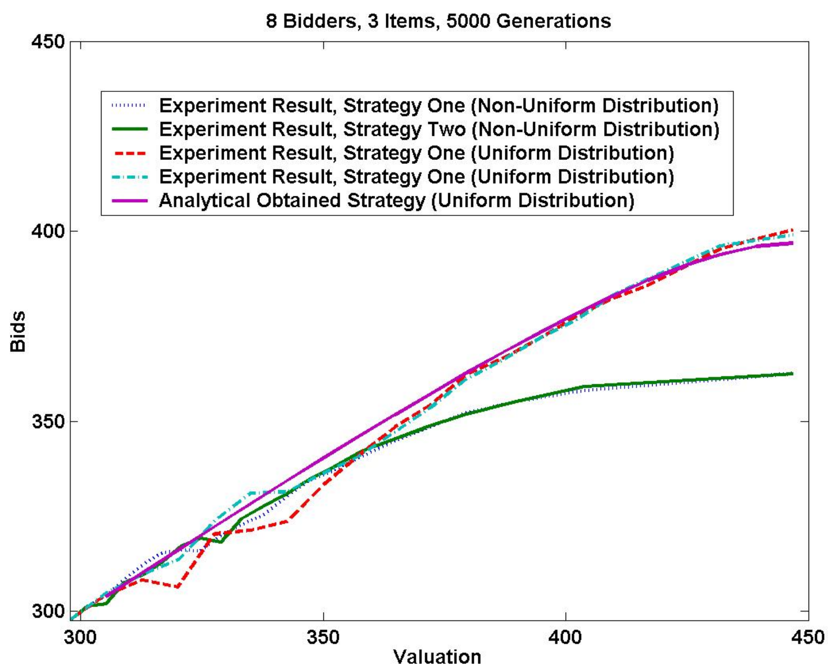

Figure 6-3: Strategy Comparison of Experiment Two

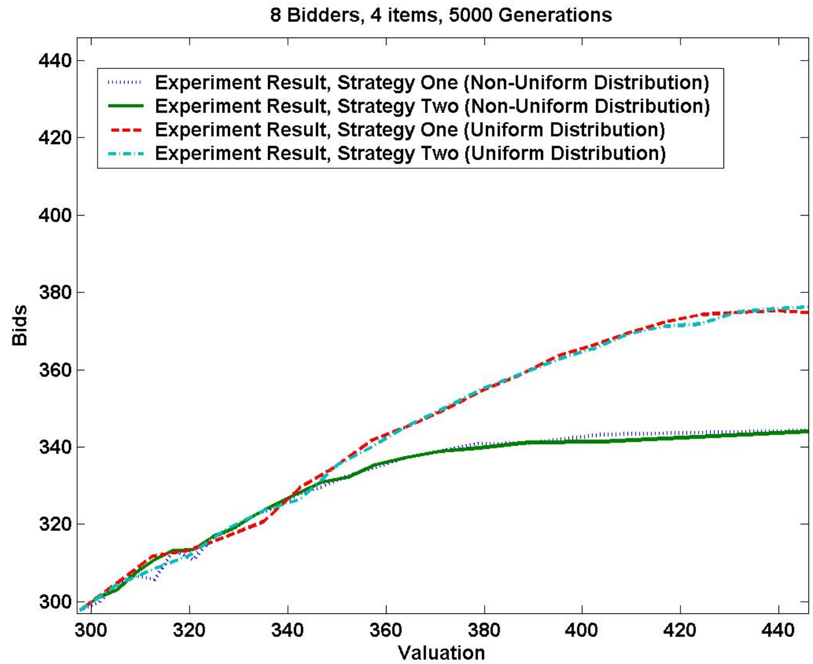

Figure 6-4: Bidding Strategies of Experiment Three

6.2.2 *Experiment Three, Four, Five, Six and Seven--n = 8* Bidders, *k = 4, 5, 6, 7* and *8* Items

In all five experiments, the seller's α chromosome converges to $\alpha = 0.9999$ and the P chromosome does not converge. This is also because the fixed price, P, is never tested; and therefore does not converge.

As observed in Chapter 4, it is not possible to obtain optimal bidding strategies using Bayesian analysis when k is large; hence, in figure 6-4, 6-5, 6-6, 6-7 and 6-8, we cannot compare the bidding strategies obtained from GA simulations to analytical results. However, it is clear from the experimental results that pure bidding strategies exist at equilibrium.

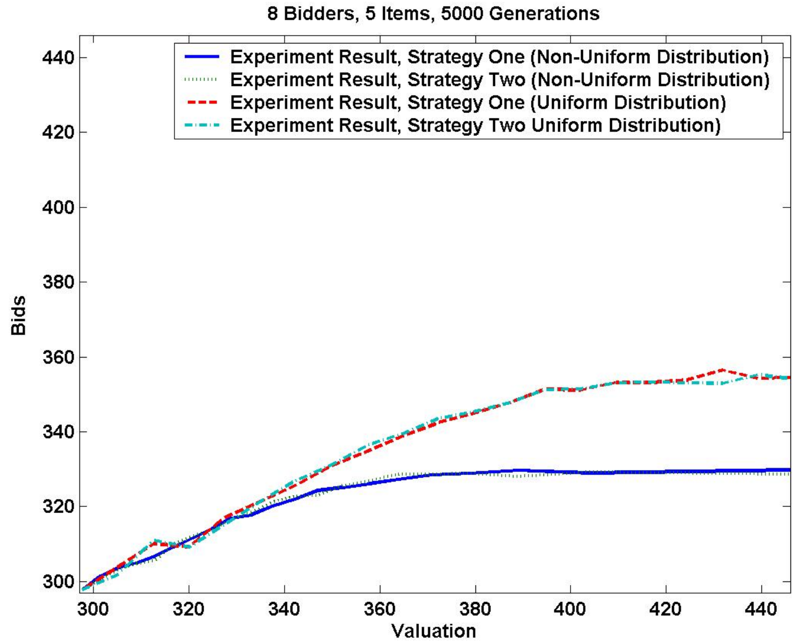

Figure 6-5: Bidding Strategies of Experiment Four

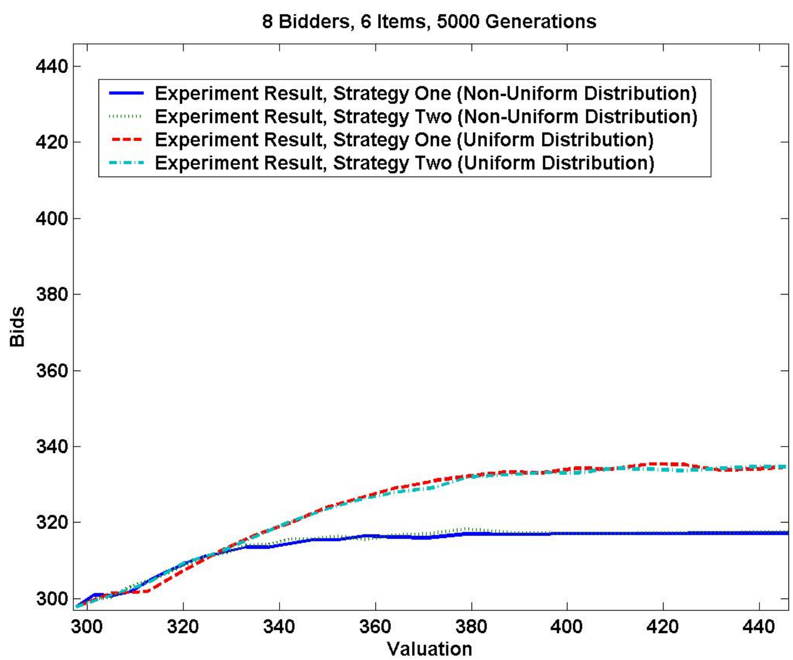

Figure 6-6: Bidding Strategies of Experiment Five

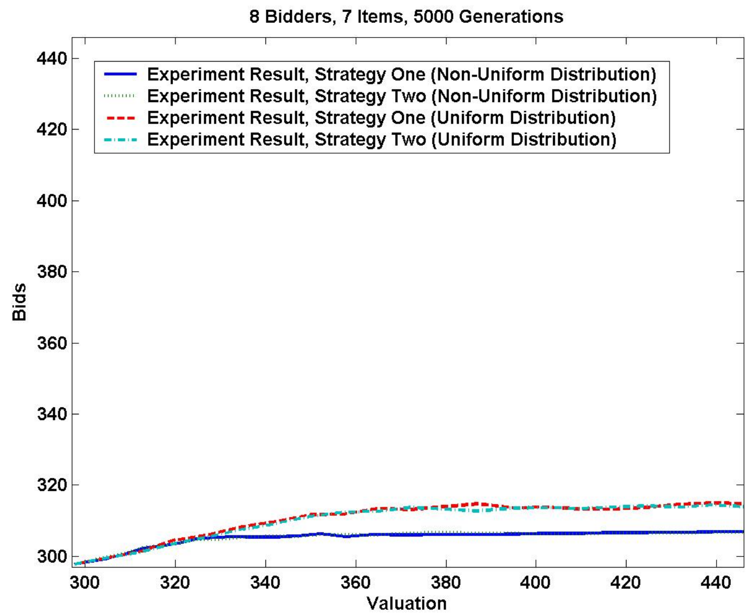

Figure 6-7: Bidding Strategies of Experiment Six

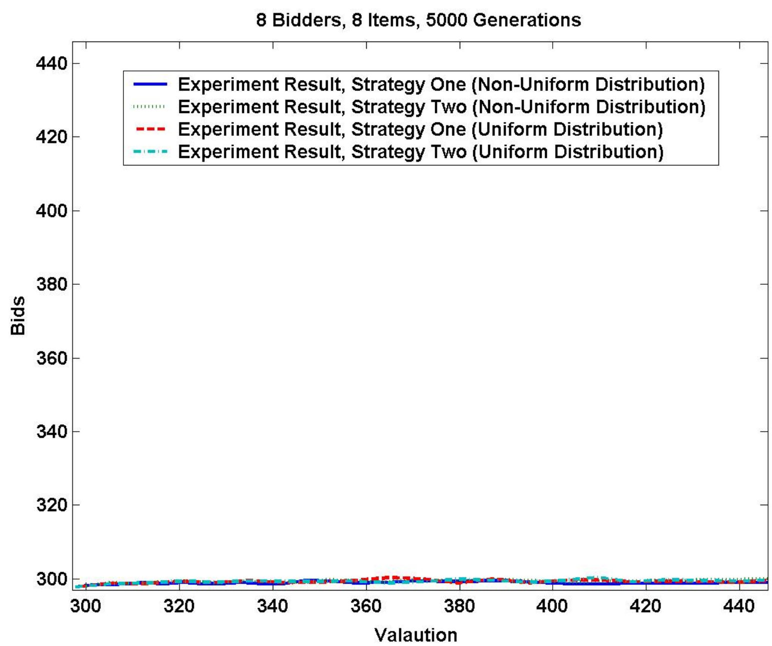

Figure 6-8: Bidding Strategies of Experiment Seven

6.2.3 *eBay Data with Optimal Bidding Strategy: n = 8 Bidders and 2 Items*

As described in section 6.1, we collected auction data from eBay with keyword "Wii 14". Data from 140 auctions were used to obtain bidders' valuation distribution over a 2 week period. The second highest bid of every auction is collected. We assume that the bid equals to valuation for all the second highest bidders because they bid as high as they can before dropping out. We take the valuation as input and compute a new optimal bid based on the bidding strategy obtained from experiment one. We assume that there are only 2 identical items available in all the 140 auctions. We calculate the difference between the actual bid submitted on eBay and our suggested, optimal bid and consider the difference to be possible saving (payoff) for all the second highest bidders. It is also the additional saving for the highest bidders. We plot the histogram of the bidder payoff if strategies obtained from experiments are adopted. The average saving for every bidder is 22.01 dollars.

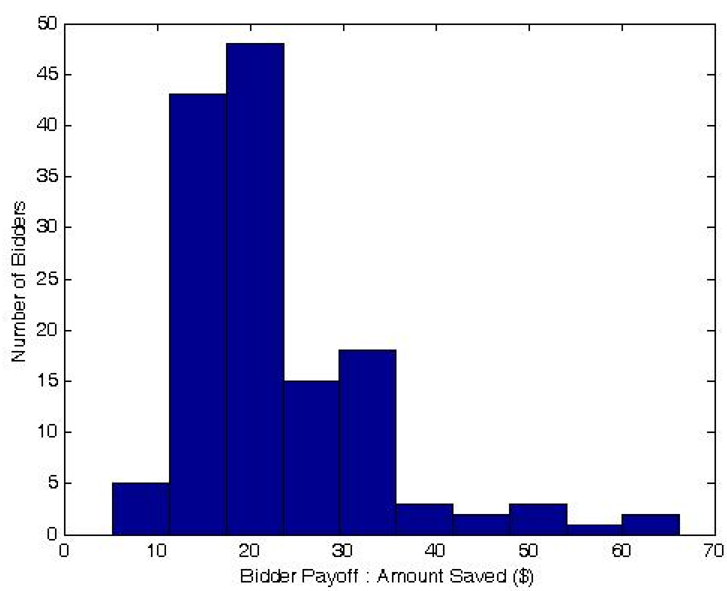

Figure 6-9: Bidder Payoff

6.3 Experiment Set C: Summary

We summarize our experimental results as follows, in this 8-bidder-2-stage game:

1. The seller's optimal strategy is to always price discriminate regardless of *k*.

   Because $\alpha = 1$, there is a corresponding pure bidding strategy for any number of items.

2. The bidder's best response is to shade (reduce) its bids according to (1) the value of *k* and (2) distribution of bidder valuation.

To summarize experiment set C, when the second-price auction is adopted at stage I, it is clear that different degrees of item scarcity make a difference for bidders' optimal strategies but not for the sellers. Sellers would always prefer privacy infringement mechanism. It could be because of the following reasons: (1) second-price auction mechanism is used at stage I, it may push bidders to bid higher even there are more than one stage, (2) the bidder valuation ranged from $297 to $446—different from experiment A with a range between 0 and 1.

6.4 Major Component and Code Flow

There are four major components of this bidding tool, including (1) an interactive user interface component, (2) an eBay data collecting component, (3) a valuation distribution computing component, (4) a genetic algorithm experiment component. Figure 6-10 illustrates the whole process, beginning with taking user input to computing a suggested bid.

6.5 Module and Programming Language Used

The entire bidding tool is written using the Python programming language. For the interactive user interface component, the EasyGUI module, an open source project written in Python, is adopted. It includes basic GUI features such as a message box, and can be programmed to accept user input. . For the eBay data collection and valuation distribution components, open source modules including easyBay, SciPy, NumPy modules are used. The easyBay module translates XML results gathered from querying eBay's API into Python objects for further computation. SciPy and NumPy are scientific and mathematic computing modules that provide features including least square fitting and polynomial roots finding. The genetic algorithm experiment component is made of reused code from previously developed experiments described in section 6.2.

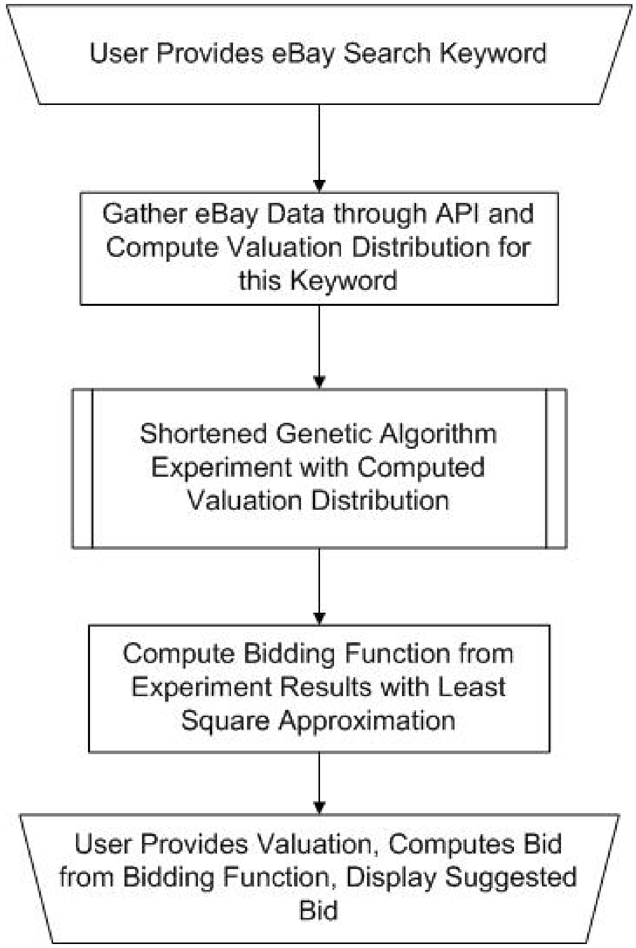

Figure 6-10 Flow Chart of the Bidding Tool

6.6 Usage Illustration and Limitation

The following screen shots are taken to illustrate the usage of this bidding tool. It first asks the user to provide keywords for the target item. After computing the bidding function, it asks the user to provide the maximum he or she is willing to bid for such an item, and then returns a suggested bid. The example shows a user searching for a Wii game console as a bundle of 14 games. The user has is willing to pay $180.50 for such an item.

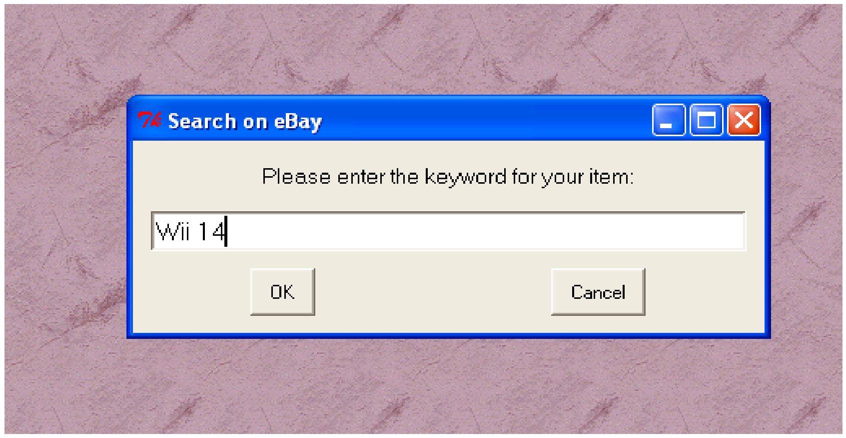

Figure 6-11: User Provides Keyword

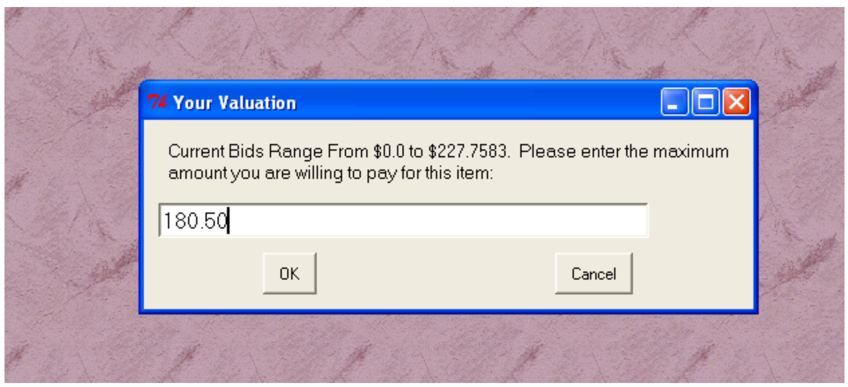

Figure 6-12: User Provides Valuation

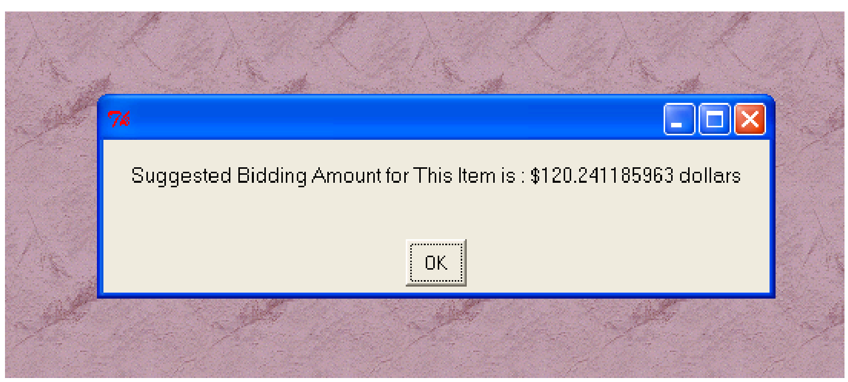

Figure 6-13: Tool Provides Bid Suggestion

A full scale genetic algorithm experiment as described in section 6.2 takes at least one week to complete on a machine of Intel Core Duo 1.8GHz processor with 2GB memory. It is not feasible for the bidding tool since auction listing on eBay has a time limit that can be as short as 3 days. To shorten the running time for the genetic algorithm experiment component, this component assumes the following conditions:

- Each generation consists of 1000 auctions
- A total number of 100 generations
- 200 sellers and each seller has 8 bidders
- Mutation rate is set to be 1/22

As a result of the shortened simulation, the bidding function has to be approximated with a least squares method. This reduces the accuracy of the optimal bidding strategy, but can be completed in 1.5 hours on a machine with Intel Core 2 Duo 2.1GHz processor and 4GB memory The number of bidders and available items can be easily modified if more features are provided for the user to specify auction to participate by identify the eBay auction ID. The bidding tool can gather information for a specific auction such as number of current bidders and take it as input for genetic algorithm experiments.

The number of available items (the value of k) is not publicly available information at this time. Because this information is not available, the bidding tool does not have a means of estimating seller reputation: the frequency with which the seller utilizes the second chance offer mechanism. However, if eBay were to make available information on seller reputation, this information could easily be used by the tool to improve the accuracy of bidding strategies.

6.7 Observation

An intelligent bidding tool can be helpful as a privacy protection mechanism. We have developed one based on our analysis of a multi-stage game that takes repeated re-encounters into consideration. It is important to note that the suggested bid does not guarantee winning the auction, as whether an auction is won depends on both, the bid and the valuations of other participating bidders. However, the suggested bid comes close to maximizing bidder payoff if the seller has one extra identical item.

The trade-off between running time and bidding strategy accuracy can be improved as computing power further evolves. Because the bidding tool is written in Python, it can be easily executed on both Windows and Linux/Unix platforms.

CHAPTER 7: CONCLUSION AND FUTURE WORK

In this dissertation, we have proposed that privacy be approached from a game-theoretic perspective. We have used the perspective to study the specific problem of the second-chance offer, by presenting a quantitative model of the second-chance offer, and a (randomized) generalization of the deterministic game of *[Joshi, Sun and Vora 05*]. We have used this approach to examine the feasibility of rationality as a privacy-protection mechanism in auctions. We are the first to study randomized seller strategies in auctions, and privacy games without closed-form solutions.

We have examined this game with both first-price and second-price auctions in stage I. Real auction data on eBay is collected as part of input to our genetic algorithm experiments. We have also examined the case of item scarcity. We have presented the results—obtained through both, Bayesian analysis and experimental results conducted with genetic algorithm simulations. It is shown that rationality provides sufficient privacy protection when items are not scarce, but not otherwise. We have also implemented bidding software as a proof of concept to utilize optimal bidding strategies on eBay. Rationality can provide limited privacy protection as automatic bidding tools. In summary, we have shown that rational behavior can perform the task of privacy protection, and that this can be implemented as a rational bidding tool in the security infrastructure. In particular, cryptographic schemes are not the only solutions for privacy protection.

We have conducted three sets of genetic algorithm experiments. In experiment set A, we are able to obtain pure bidding strategies for $n = 8$ bidders and $k = 2, 3, 4, 5$ items. We have observed the existence of possible mixed strategies for $n = 8$ bidders and $k = 6, 7, 8$ items. We further conduct experiment set B and successfully obtain the mixed strategies by changing the encoding of bidder chromosomes. In experiment set C, we apply the techniques developed in previous experiment sets with real eBay data. We are able to obtain pure bidding strategies when second-price auction mechanism is used in stage I.

Possible future work includes (1) generalizing the privacy model of a two-stage game into a repeated, infinite stage game, (2) applying genetic algorithm experiment for other problems with interdependent objective functions, and (3) applying the encoding scheme developed in this dissertation to obtain mixed strategies in other type of games and exploring other encoding schemes as well.

REFERENCES

[Ackerman, Cranor and Reagle 99] Mark S. Ackerman, Lorrie F. Cranor, and Joseph Reagle, “Privacy in e-Commerce: Examining User Scenarios and Privacy Preferences”, In Proceedings of the ACM Conference on Electronic Conference (EC’99), November 1999

[Acquisti 04] Aleesandro Acquisti, “Privacy in Electronic Commerce and the Economics of Immediate Gratification”, In *Proceedings of the ACM Conference on Electronic Conference (EC’04),* May 2004

[Acquisti, Dingleine and Syverson 03] Alessandro Acquisti, Roger Dingledine, and Paul Syverson, “On the Economics of Anonymity”, In *Financial Cryptography (FC 03),* Springer Verlag, LNCS, 2003

[Acquisti and Grossklags 04] Alessandro Acquisti and Jens Grossklags, “Privacy and Rationality: Preliminary Evidence from Survey Data”, In *Third Workshop on Economics and Information Security,* University of Minnesota, Minneapolis, MN, May 2004

[Acquisti and Varian 01] Acquisti and Hal R. Varian, “Conditioning Prices on Purchase History”, *Technical Report*, Unibersity of California, Berkeley, 2001

[Andreoni and Miller 95] James Andreoni and John H. Miller, “Auctions with Artificial Adaptive Agents”, *Games and Economic Behavior*, 10:3964, 1995

[Arifovic 94] Jasmina Arifovic, “Genetic Algorithm Learning and the Cobweb Model”, *Journal of Economic Dynamics and Control*, vol. 18(1), 1994

[Axelrod 84] Robert Axelrod, “Evolution of Cooperation”, Basic Books, 1984

[Axelrod 87] Robert Axelrod, "The Evolution of Strategies in the Iterated Prisoner's Dilemma", Morgan Kaufman, 1987

[Byde 03] Andrew Byde, "Applying Evolutionary Game Theory to Auction Mechanism Design", In *Proceedings of the IEEE International Conference on Electronic Commerce (CEC03),* May 2003

[Chong and Zack 96] Edwin K. P. Chong and Stanislaw H. Zak, "An Introduction to Optimization", Wiley Press, 1996

[Cliff 03] Dave Cliff, "Explorations in Evolutionary Design of Online Auction Market Mechanisms", Journal *of Electronic Commerce Research and Applications,* vol. 2(2), 2003

[Cliff 06] Dave Cliff, "ZIP60: Further Explorations in the Evolutionary Design of Trader Agents and Online Auction-Market Mechanisms", *IEEE Transactions on Evolutionary Computation,* 2006

[Conitzer and Sandholm 03] Vincent Conitzer and Tuomas Sandholm, "Complexity Results about Nash Equilibria", In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI-03),* pp. 765-771, Acapulco, Mexico, 2003

[Crawford and Sobel 82] Vicent P. Crawford and Joel Sobel, "Strategic Information Transmission", *Econometrica*, Econometric Society, vol. 50(6), 1982

[Dawid 96] Herbert Dawid, "Genetic Algorithms as a model of adaptive learning in economic systems", *Central European Journal for Operations Research and Economics* 4(1): 7-23, 1996

[Dawid 99] Herbert Dawid, "Adaptive Learning by Genetic Algorithms: Analytical Results and Applications to Economic Models", Springer, 1999

[Dawid and Kopel 98] Herbert Dawid and Michael Kopel, "On Economic Applications of the Genetic Algorithm: a Model of the Cobweb Type", *Journal of Evolutionary Economics* 8(3), 1998

[Dorigo 92] Marco Dorigo, "Optimization, Learning and Natural Algorithms", *PhD thesis*, Politecnico di Milano, Italy, 1992

[Goldberg 88] David E. Goldberg, "Genetic Algorithms in Search, Optimization, and Machine Learning", Addison-Wesley, 1988

[Haupt and Haupt 05] Randy L. Haupt and Sue Ellen Haupt, "Practical Genetic Algorithms", 2nd Edition, Wiley Press, 2004

[Holland 75] John H. Holland, "Adaptation in Natural and Artificial Systems", Ann Arbor: Univ. of Michigan Press, 1975 (Also see 2nd edition, MIT Press, 1992)

[Holland 92] John H. Holland, "Genetic Algorithms", *Scientific American*, 66-72, July 2004

[Joshi, Sun and Vora 05] Sumit Joshi, Yu-An Sun and Poorvi L. Vora, "The Privacy Cost of the Second Chance Offer", *ACM Workshop on Privacy in the Electronic Society (WPES 05),* 2005

[Joshi, Sun and Vora 06] Sumit Joshi, Yu-An Sun and Poorvi L. Vora, "Randomization as a Seller Strategy during Price Discrimination, and its Impact on Bidder Privacy", *ACM Worshop on Privacy in the Electronic Society (WPES 06),* 2006

[Joshi, Sun and Vora 07] Sumit Joshi, Yu-An Sun and Poorvi L. Vora, “Privacy and Private Auctions”, *Working Paper*, 2007

[Klemperer 04] Paul Klemperer, “Auctions: Theory and Practice”, Princeton University Press, 2004

[Krishna 02] Vijay Krishna, “Auction Theory”, Academic Press, 2002

[Laudon 96] Kenneth Laudon, “Market and Privacy”, *Communication of the ACM*, 1996

[Lucking-Reiley 00] David Lucking-Reiley, “Vickrey Auctions in Practice: From Nineteenth Century Philately to Twenty-First Century E-Commerce”, *Journal of Economics Perspectives*, vol. 14(3), 183-192, 2000

[McAfee and McMillan 87] R. Preston McAfee and John McMillan, “Auctions and Bidding”, *Journal of Economic Literature,* American Economic Association, vol. 25(2), 1987

[Menezes and Monteiro 05] Flavio M. Menezes, Paulo K. Monteiro, “Introduction to Auction Theory”, Oxford University Press, 2005

[Milgrom 89] Paul Milgrom, “Auctions and Bidding: A Primer”, *Journal of Economic Perspectives*, vol. 3(3), 1989

[Milgrom and Weber 82] Paul R. Milgrom and Robert J. Weber, “A Theory of Auctions and Competitive Bidding”, *Econometrica*, Econometric Society, vol. 50(5), 1982

[Mühlenbein and Schlierkamp-Voosen 93] H. Mühlenbein and D. Schlierkamp-Voosen, “Predictive Models for the Breeder Genetic Algorithm: I. Continuous Parameter Optimization”, *Evolutionary Computation*, 1 (1), pp. 25-49, 1993

[Odlyzko 03] Andrew Odlyzko, "Privacy, Economics and Price Discrimination on the Internet", ICEC 2003, In *Proceedings of the Fifth International Conference on Electronic Commerce*, ACM Press 2003

[Pindyck and Rubinfeld 04] Robert S. Pindyck and Daniel L. Rubinfeld, "Microeconomics", 6th Edition, Prentice Hall, 2004

[Riechmann 01] Thomas Riechmann, "Learning in Economics: Analysis and Application of Genetic Algorithms", Physica-Verlag Heidelberg, 2001

[Rothkopf, Teisberg and Kahn 90] Michael H. Rothkopf, Thomas J. Teisberg and Edward P. Kahn, "Why Are Vickrey Auctions Rare?", *Journal of Political Economy*, vol. 98(1), 1990

[Salmon and Wilson 06] Timothy C. Salmon and Bart J. Wilson, "Second Chance Offers versus Sequential Auctions: Theory and Behavior", *Economic Theory* 2006

[Varian 00] Hal R. Varian, "Economic Aspects of Personal Privacy", In Brian Kahin and Hal R. Varian, editors, *Internet Publishing and Beyond*, MIT Press 2000

[Vickrey 61] William Vickrey, "Counterspeculation, Auctions and Competitive Sealed Tenders", *Journal of Finances*, vol. 16(1), 1961

[Walia, Byde and Cliff 03] Vibhu Walia, Andrew Byde and Dave Cliff, "Evolving Market Design in Zero-Intelligence Trader Markets", In *Proceedings of the IEEE International Conference on Electronic Commerce (CEC03)*, May 2003

APPENDICES

Appendix A

Proof: Theorem 2. The expected payoff due to bid *b*, when all others are bidding according to strategy *β*, is, from section 4.1:

$$E[\Pi(b,x)] = \begin{cases} G(\beta^{-1}(b))(x-b) + H(\beta^{-1}(b))[\alpha(x-b) + (1-\alpha)(x-P)] & x > P \\ G(\beta^{-1}(b))(x-b) + H(\beta^{-1}(b))[\alpha(x-b)] & else \end{cases} \quad (1)$$

To find the bidding strategy *β** that maximizes the expected payoff, we differentiate (1) wrt *b*, equate to zero, and, assuming a symmetric strategy among bidders, replace *b* with *β*(x)*. This gives us:

$$\alpha h(x)\beta^*(x) + \alpha\beta^{*\prime}(x)H(x) + \beta^*(x)g(x) + G(x)\beta^{*\prime}(x)$$
$$= \begin{cases} xg(x) + \alpha xh(x) + (1-\alpha)(x-P)h(x) & x > P \\ xg(x) + \alpha xh(x) & else \end{cases}$$

Integrating both sides wrt *x* gives:

$$\alpha\beta^*(x)H(x) + \beta^*(x)G(x)$$
$$= \begin{cases} \int_0^x yg(y)dy + \alpha\int_0^x yh(y)dy + (1-\alpha)\int_x^P h(y)(y-P)dy & x > P \\ \int_0^x yg(y)dy + \alpha\int_0^x yh(y)dy & else \end{cases}$$

Solving for *β** gives:

$$\beta^*(x) = \begin{cases} \dfrac{\int_0^x yG'(y)dy + \alpha\int_0^x yH'(y)dy + (1-\alpha)\int_P^x (y-P)H'(y)dy}{G(x)+\alpha H(x)} & x > P \\ \dfrac{\int_0^x yG'(y)dy + \alpha\int_0^x yH'(y)dy}{G(x)+\alpha H(x)} & else \end{cases} \quad (2)$$

When x is uniformly distributed, $G(x) = x^{n-1}$; and

$$\frac{\int xG'(x)dx}{G(x)} = \frac{n-1}{n}x$$

may be substituted in (2).

We now show that unilateral deviation by a single bidder does not provide benefit. Clearly, there is no value in submitting a bid greater than the highest value of $\beta^*(x)$ because the bidder can certainly win with a bid that is equal to this highest value. Similarly, there is no value in submitting a bid lower than the lowest value of $\beta(x)$, as this bid will certainly not be a winning bid. Hence a deviating bidder will only provide a bid from the range of β^*.

Suppose a bidder with valuation x bid $\beta^*(z)$, $z \neq x$. We consider cases when both x and z are smaller than and greater than P, and, also when $x < P < z$ and $z < P < x$.

Consider the case when $x < P$ and $z < P$. The bidder's payoff is:

$$E[\Pi(\beta^*(z), x)] = G(z)(x - \beta^*(z)) + H(z)[\alpha(x - \beta^*(z))] \tag{3}$$

Further,

$$E[\Pi(\beta^*(x), x)] = G(x)(x - \beta^*(x)) + H(x)[\alpha(x - \beta^*(x))] \tag{4}$$

To show that bidders do not have an incentive to unilaterally deviate, we need to show that $E[\Pi(\beta^*(x), x)] - E[\Pi(\beta^*(z), x)] \geq 0$ for $x \leq z$ and for $x > z$.

Substituting (2) in (3) and (4) gives:

$$E[\Pi(\beta^*(x), x)] = \int_0^x G(y)dy + \alpha\int_0^x H(y)dy \tag{5}$$

and:

$$E[\Pi(\beta^*(z),x)] = G(z)(x-z) + \alpha H(z)(x-z) + \int_0^z G(y)dy + \alpha \int_0^z H(y)dy \qquad (6)$$

Subtracting (6) from (5) gives:

$$\begin{aligned} & E[\Pi(\beta^*(x),x)] - E[\Pi(\beta^*(z),x)] \\ & = (z-x)[G(x) + \alpha H(x)] - \int_x^z [G(y) + \alpha H(y)]dy \\ & = (z-x)[(1-\alpha)G(z) + \alpha K(z)] - \int_x^z [(1-\alpha)G(y) + \alpha K(y)]dy \end{aligned}$$

where *K(x)* is the probability that x is among the k highest bidders, *K(x) = G(x) + H(x)*. The expression above is non-negative, by an argument similar to that provided for the derivation of the equilibrium bidding strategy in a classical first-price auction in [Krishna 02], for both $x \geq z$ and $x \leq z$, because both *G(x)* and *K(x)* are monotonic increasing.

We now examine the case for $x \geq P$ and $z \geq P$. The expected payoff for bidding *β*(z),* when $x \geq P$ is

$$E[\Pi(\beta^*(z),x)] = G(z)(x-\beta^*(z)) + H(z)[\alpha(x-\beta^*(z)) + (1-\alpha)(x-P)] \qquad (7)$$

Further,

$$E[\Pi(\beta^*(x),x)] = G(x)(x-\beta^*(x)) + H(x)[\alpha(x-\beta^*(x)) + (1-\alpha)(x-P)] \qquad (8)$$

Again, we need to show that $E[\Pi(\beta^*(x),x)] - E[\Pi(\beta^*(z),x)] \geq 0$, whether $x \leq z$ and for $x > z$. Again, by substituting the expression for *β*(x),* we obtain:

$$
\begin{aligned}
&E[\Pi(\beta^*(x),x)] \\
&= \int_0^x G(y)dy + \alpha \int_0^x H(y)dy + (1-\alpha)[(x-P)H(x) - \int_P^x (y-P)h(y)dy] \qquad (9) \\
&= \int_0^x G(y)dy + \alpha \int_0^x H(y)dy + (1-\alpha)\int_P^x H(y)dy
\end{aligned}
$$

Similarly,

$$
\begin{aligned}
&E[\Pi(\beta^*(z),x)] \\
&= G(z)(x-z) + \alpha H(z)(x-z) + \int_0^z G(y)dy + \alpha \int_0^z H(y)dy + (1-\alpha)[(z-P)H(z) \\
&- \int_P^z (y-P)h(y)dy]
\end{aligned}
$$

or:

$$
\begin{aligned}
&E[\Pi(\beta^*(z),x)] \\
&= K(z)(x-z) + \int_0^z G(y)dy + \alpha \int_0^z H(y)dy + (1-\alpha)\int_P^z H(y)dy \qquad (10)
\end{aligned}
$$

Hence

$$
E[\Pi(\beta^*(x),x)] - E[\Pi(\beta^*(z),x)] = (z-x)K(z) - \int_x^z K(y)dy \qquad (11)
$$

Again, the above expression is non-negative because *K(x)* is monotonic increasing.

Now we examine the case when $x < P < z$. Using (10) and (5) we obtain:

$$
\begin{aligned}
&E[\Pi(\beta^*(x),x)] - E[\Pi(\beta^*(z),x)] \\
&= (z-x)[G(z) + \alpha H(z)] + \int_z^x [G(y) + \alpha H(y)]dy - (1-\alpha)\int_P^z H(y)dy + (z-x)(1-\alpha)H(z) \geq 0
\end{aligned}
$$

by argument as with the case $z \leq P$ and $x \leq P$.

Thus we have shown that bidders do not have an incentive to deviate from β^* when it is the strategy used by all other bidders. Hence β^* is an equilibrium strategy if $0 \leq \beta^* \leq x$ and β^* is monotonic increasing.

Appendix B

From section 4.3, we know that the expected payoff for the two-stage game with second-price auction is:

$$E[\prod] = G(\beta^{-1}(b))\left(x - \frac{\int_0^x bG'(y)dy}{G(x)}\right) + H(\beta^{-1}(b))[\alpha(x-b)] \quad \text{where } x < P \tag{1}$$

To find the bidding strategy β^* that maximizes the expected payoff, we differentiate (1) wrt *b*, equate to zero, and, assuming a symmetric strategy among bidders, replace *b* with $\beta^*(x)$. This gives us:

$$xg(x) - \beta(x)g(x) + \alpha xh(x) - \alpha H(x)\beta'(x) - \alpha\beta(x)h(x) = 0 \tag{2}$$

To solve the partial differential equation, we have to substitute $H(x)$ for specific n, k and α. In the following examples, we substitute $\alpha = 1$ based on experiment results obtained in chapter 6. We also assume x is uniformly distributed.

<u>For $n = 8$ bidders and $k = 2$ items:</u>

We have $H(x) = 7x^6 - 7x^7$, $G(x) = x^7$, and $h(x) = H'(x) = 42x^5 - 49x^6$. Substituting (2) with $H(x)$, $G(x)$ and $h(x)$, it gives:

$$\beta^{*\prime}(x) = \frac{6x - 6\beta^*(x)}{x} \tag{3}$$

Using Matlab, we can obtain the solution of (3) to be:

$$\beta^*(x) = \frac{6}{7}x$$

For $n = 8$ bidders and $k = 3$ items:

We have $H(x) = 14x^7 - 35x^6 + 21x^5$, $G(x) = x^7$, and $h(x) = 98x^6 - 210x^5 + 105x^4$.

Similarly, substituting (2) with *H(x), G(x)* and *h(x),* it gives:

$$15x^7 - 30x^6 + 15x^5 = \beta^*(x)(15x^6 - 30x^5 + 15x^4) + \beta^{*\prime}(x)(2x^7 - 5x^6 + 30x^5) \qquad (4)$$

Using Matlab, we can obtain the solution of (4) to be:

$$\beta^*(x) = \frac{1}{17} + \frac{3}{17017}\left(\frac{-5832 + 45x^4 + 324x^2 - 1944i(6x-9)^{1/2} + 1944x + 9030x^5}{x^5(2x-3)^3} + \right.$$
$$\left. \frac{+108x^3 + 145635x^6 - 180180x^8 + 40040x^9}{x^5(2x-3)^3}\right)$$