Multi-Stage Information Acquisition in Auction Design

A thesis presented
by

Kyna G. Fong

To
Applied Mathematics
in partial fulfillment of the honors requirements
for the degree of
Bachelor of Arts
Harvard College
Cambridge, Massachusetts

April 1, 2003

Contents

Abstract

In traditional auction theory, bidders in auction mechanisms are assumed to know their exact values for the goods being sold. Under such assumptions, auction theorists have arrived at the *revelation principle*, which implies that the English ascending-price auction yields the equivalent outcome as the Vickrey second-price sealed-bid auction. In reality, however, a bidder's private value for a good is often uncertain. In this paper, we apply concepts of bounded rationality from artificial intelligence to traditional auction theory to arrive at the conclusion that the revelation principle no longer holds when considering bounded-rational bidders with uncertainty in their true valuations of goods. We begin by setting up a mathematical framework for such *uninformed* bidders and derive the optimal bidding and information acquisition strategies in three canonical auction formats: the posted-price, the second-price sealed-bid and the ascending-price. Using these strategies, we find that the ascending-price auction provides much better incentives for the right bidders to acquire the right amounts of information and may lead to higher revenues than the second-price sealed-bid auction. The main achievement of this work is to introduce a consistent mathematical valuation model of mid-auction multi-stage information acquisition. Most of the previous work on information acquisition has focused acquiring before auctions begin, while some more recent work has looked at mid-auction information acquisition but considering only one stage of information acquisition. We believe that ours is the first to consider multiple stages of information acquisition in auctions with bounded-rational bidders.

Chapter 1

Introduction

In traditional auction theory, bidders in auction mechanisms are assumed to know their exact values for the goods being sold. Under such assumptions, auction theorists have arrived at the *revelation principle*, which states that for every iterative auction mechanism with multiple rounds, there exists a corresponding direct or sealed-bid mechanism, in which bidders state their values for all items at the beginning of the auction, that has the same outcome. Thus, the equilibrium strategies in such mechanisms are easily computable. Furthermore, it follows that the English ascending-price auction is equivalent to the Vickrey second-price sealed-bid auction.

In reality, however, a bidder's private value for a good is often uncertain. In fact, people can rarely name the exact value an item has for them and instead may only know an estimate or range of possible values. Moreover, obtaining a more accurate valuation may require expending significant efforts and resources. This occurs when bidders have *hard valuation problems.* The idea of bidders having hard valuation problems arises from considering the effect of bidders' *bounded rationality* in auctions. The idea of bounded rationality is a central concept in the field of artificial intelligence and describes agents that are faced with bounded computational resources. In particular, they are computationally restricted from immediately calculating perfectly rational strategies and must instead decide how to expend their resources in acquiring additional information. Such agents are closer in line with reality than perfectly rational agents.

In this paper, we apply concepts of bounded rationality from artificial intelligence to traditional auction mechanisms in economics. This is only a small facet of the fascinating interface between computer science and economics that is quickly rising in popularity, and our paper is intended to emphasize the added benefits to combining insights from both fields.

A classic example of an auction setting in which agents are faced with a costly and difficult valuation problem is that of a corporate takeover auction. In such a mechanism, the potential acquirors may spend millions of dollars determining the value of the assets for sale and calculating potential synergies. Another example is that of wireless spectrum auctions. Telecommunication companies must incur large costs in order to determine not only on which bands of spectrum they wish to bid but also how much they are willing to bid. Only significant amounts of costly research on the profitability of different options will allow these companies to refine the range of their valuations.

In practice, uncertainty about valuations in auctions arises in many ways. In complex environments such as combinatorial auctions, gathering information and computing valuations for all goods and combinations of goods can be formidable. It makes sense for bidders to approach such auctions with just a vague estimate of their valuations, and, as the mechanism progresses, then to choose to concentrate her computational resources on acquiring information on the most promising alternatives. It does not make sense for a bidder to waste her computation resources refining her value when she is not even sure that she has a decent chance at winning the good. Only when she receives feedback that she is competitive in the auction will she then invest resources in acquiring information to refine her value.

In this paper, we set up a mathematical framework in which we have bounded-rational bidders that are uncertain about their true values for a good but have *the option of acquiring information multiple times* in order to refine their beliefs about those values. We focus on *private-value settings* in which the value to each bidder of the good for sale is independent of the value to every other bidder. With this framework in mind, we consider several different auction mechanisms and determine the optimal strategies given the bounded-rationality of bidders. Armed with those optimal strategies, we can then draw conclusions about the relative merits of the different auction formats when bidders are bounded-rational. Our ultimate goal is to strive towards the optimal design of auction mechanisms for bidders with hard valuation problems.

Ideally, the optimal auction will result in as little wasted computation as possible. Bidders acquire just enough information to determine their relative competitiveness in the auction and then stop because such computation is expensive. At the same time, we also hope that the optimal auction will result in high efficiency and revenues.

Related Work

Information acquisition *during* auctions is an issue that has been ignored in auction design until relatively recently. Previously, most of the literature on information acquisition has restricted its attention to static or sealed-bid auctions in which information acquisition may take place only prior to the auction (e.g. [10] and [9]).

There are a few works of which we are aware that do deal with information acquisition *during* auctions and the three most relevant include Compte & Jehiel [4], Rezende [14], and Parkes [11, 12]. The works by Compte & Jehiel and Rezende both propose stylized equilibrium models for costly information acquisition by bidders in sealed-bid and ascending-price auctions. Bidders incur the cost of information acquisition *during* auctions and can choose to acquire information only *once* during the auction. While Rezende focuses on a model in which multiple bidders are uninformed, Compte & Jehiel provide a simpler model in which only one bidder has an uncertain valuation. These stylized models are interesting because they allow for tractable equilibrium analysis of optimal strategies. This paper serves to extend these models by providing a framework in which bidders can acquire information *multiple times* during an auction. By doing so, we are able to support the major conclusions drawn in these papers as well as explain some more subtle effects.

The paper by Parkes [11] is closest in line with the work done in this paper. Parkes considers a simple valuation model for agents by assuming that their values are uniformly distributed along some interval and then proposes an algorithm for information acquisition in this model. Unfortunately, in Parkes' model, information acquisition is inconsistent with the valuation model. Nonetheless, he does demonstrate that iterative auctions are preferred over sealed-bid auctions when bidders' valuations are uncertain. In particular, he emphasizes that iterative auctions allow for the *right* bidders to acquire the *right* amount of information.

Two other works relating to mid-auction information acquisition are also worth noting. First, Engelbrecht-Wiggans [5] has written a paper in which he solves an analytic example using two stages of information acquisition and suggests that information acquisition is the main reason for the popularity of multi-stage over single-stage auctions. Secondly, Rasmusen [13] looks at a simple model where agents can acquire information once in iterative auctions but his model is quite restrictive in that there are only two bidders, each with a very specific distribution for the value of the good.

Finally, Sandholm [17] deals with the question of mechanism design with bounded-rational agents. This analysis shows that an agent can make a better decision about whether to perform further computation about its own value if it is well-informed about other agents' bids. Larson & Sandholm [7] [8] model a full equilibrium where agents must make explicit decisions about whether to deliberate about their own values or the values of other agents (counterspeculation) in settings with costly deliberation.

We believe that our paper is the first to introduce a consistent valuation model of multi-stage information acquisition. Using that model, we arrive at optimal bidding and information acquisition strategies for bidders in three canonical auction formats: posted-price, second-price sealed-bid, and ascending-price. From our analysis, we confirm the findings of previous work and find that the ascending-price auction provides much better incentives for the right bidders to acquire the right amount of information than the second-price sealed-bid auction. The advantage is particularly obvious when bidders have no prior knowledge about the distribution of the other bidders' values. We also find that this advantage may lead to higher revenues in the ascending format when multiple bidders have hard valuation problems.

Motivational Examples

In order to further motivate our analysis, we describe some basic examples of situations in which we may find hard valuation problems. By laying out these examples, we can focus on what exact questions we need to answer in our analysis.

First, imagine a man from Boston trying to purchase a last-minute ticket online to San Francisco. His favorite discount online airfare finder quotes him a price of $500. But he is unsure of how much the trip is worth to him, what exact value a ticket to San Francisco gives him. Should he spend time finding out more details about his trip? Even if he chooses to find out more details, should he accept the price?

Now imagine a young couple looking for a new home. Assume that to buy a specific home, the couple, along with any other parties interested in the home, must submit a bid to the home's owner specifying the amount they are willing to pay. One day the couple spots a home that seems to fit their needs almost perfectly, so they decide to submit a bid. But they do not know all the details and intricacies of the house's characteristics such as its design, structure, insulation, school district, location, etc. How much time should the couple spend researching information about the house? What bid should they finally submit?

Lastly consider an avid art collector participating in a Sotheby's online auction. A certain painting has caught his interest, and he decides to start bidding for it. He realizes, however, that he is not sure whether the painting is genuine, what it looks like in reality, and how it will fit into his collection. How much of this uncertainty should he try to resolve before bidding? At what value should he give up?

In each of these examples, the bidder is participating in a different auction format (the posted-price, second-price sealed-bid, and ascending-price auctions, respectively) but still faces a common decision problem. Solving that decision problem for each of the different auction formats is where our interest lies.

Chapter 2

Setup

In this chapter, we lay out the setup for our analysis. First, we formalize the bidder's decision problem and establish the analytical frameworks with which we model that problem. Next, we specify the auction formats that we have chosen to analyze. Finally, we outline some stylized numerical examples to motivate our analysis.

2.1 The Three-Part Decision Problem

The bidder's decision problem can be divided into three parts: (i) the valuation problem, (ii) the bidding problem, and (iii) the information acquisition problem. We describe each of these parts in detail below.

First, we list a few basic assumptions. We have a set of n bidders, indexed by $i \in \{1, 2, ..., n\}$. Each bidder i has a *private* valuation function $v_i : G \to \mathbb{R}^+$ for goods in a set G, and a von-Neumann Morgenstern utility function for money $U_i : \mathbb{R} \to \mathbb{R}$, which for simplicity is normalized to $U_i(0) = 0$. So, the utility to a bidder i of buying a good g at price p is $U_i(v_i(g) - p) = v_i(g) - p$. Also, as is standard in auction theory, we assume risk-neutral bidders with linear-separable and linear-additive utility functions. Each bidder has an independent private valuation for goods. Since we only consider markets for a single good g, we represent $v_i(g)$ as simply v_i. In laying out the basic model for the three-part decision problem below, we assume that there is only a single good for sale.

2.1.1 Modeling the Valuation Problem

The bidder is faced with the *valuation problem* when she needs to figure out how much she values the good for sale. In this paper we want to deal with the oft-occurring instances in which a bidder's valuation problem is hard. In such cases, the bidder has uncertainty about her exact private value for the good being auctioned and must incur costs to acquire information and thereby to refine her beliefs. In this section, we develop a simple theoretical model that supports multiple stages of information acquisition. To do so, we employ the use of the *Gaussian*, or *normal*, distribution. This model has been proposed by Chiburis [3].

Consider bidder i with a hard valuation problem and thus uncertainty about her value. We call such a bidder *uninformed*. We model her beliefs about her true value as a Gaussian distribution, represented by the function f_{μ^i,σ^i}, where μ^i is the mean and σ^i is the standard deviation. So, our family of valuation distributions is $f_{\mu,\sigma}(\cdot)$ where f is a Gaussian distribution with mean μ and standard deviation σ. In our analysis, uninformed bidders have the option to *acquire information*. We model one step of information acquisition in the following way and show that it is *consistent* with our model for the bidder's beliefs.[1]

Proposition 1. *Let $0 < \alpha < 1$ be a constant factor such that $1 - \alpha$ represents the computational effectiveness of one information acquisition step. Suppose a bidder has a Gaussian valuation distribution $f_{\mu,\sigma}$ about her true value v for some good. It is then consistent with that belief distribution to perform one step of information acquisition by updating $f_{\mu,\sigma}$ to a new valuation distribution $f_{\mu',\sigma'}$, where (i) $f_{\mu',\sigma'}$ is also Gaussian, (ii) $\sigma' = \alpha\sigma$, and (iii) μ' is chosen according to a Gaussian density with mean μ and standard deviation $\sqrt{1-\alpha^2}\sigma$.*

Proof. Consider a bidder with a valuation distribution $f_{\mu,\sigma}$. Let $h_{\mu,\rho}$ be the function that chooses the mean of the new distribution after acquiring information. In other words, if a bidder's valuation distribution before acquiring information is $f_{\mu,\sigma}$, the new mean y after acquiring information is chosen according to the distribution $h_{\mu,\rho}$. We assume that the standard deviation of the new distribution decreases by a factor of α (where $1 - \alpha$ is a measure of the computational effectiveness).

For any value x, we want $f_{\mu,\sigma}(x)$, i.e. the *ex ante* probability distribution of x, to be equal to the expected *ex post* distribution of x, which is the integral over all possible new

[1] It is important that the model for one step of information acquisition be consistent with the valuation model so that the model can be extended to many steps of information acquisition. As an example, Parkes [11] models agents with uniform distributions and we find that in such a model information acquisition is inconsistent with the agents' beliefs! As a result, the extension of Parkes' model to many steps of value refinement is not viable.

distributions of the probability of x in the new distribution. Symbolically, we have that

$$\begin{aligned} f_{\mu,\sigma}(x) &= \int_{-\infty}^{\infty} \Pr[y \text{ is new mean}] \Pr[x|y \text{ is new mean}] \, dy \\ &= \int_{-\infty}^{\infty} h_{\mu,\rho}(y) f_{y,\alpha\sigma}(x) \, dy. \end{aligned}$$

That integral is known as the *convolution* of $h_{\mu,\rho}(y)$ with $f_{y,\alpha\sigma}(x)$. We need to prove that $h_{\mu,\rho}$ can in fact be a Gaussian density with mean μ and standard deviation $\sqrt{1-\alpha^2}\sigma$. In other words, we need to prove that

$$f_{\mu,\sigma}(x) = \int_{-\infty}^{\infty} f_{\mu,\sqrt{1-\alpha^2}\sigma}(y) f_{y,\alpha\sigma}(x) \, dy \tag{2.1.1}$$

where $0 < \alpha < 1$. By plugging in the explicit function for a Gaussian into the above formula, Chiburis has shown that the above equation holds true [3] (see Appendix for proof). Thus, we conclude that our model for one step of information acquisition is consistent. □

Because of the consistency of this Gaussian model of information acquisition and belief updating, we can extend the result to $m \geq 1$ *consecutive* steps of information acquisition.

Proposition 2. *Let $0 < \alpha < 1$ be a constant factor such that $1 - \alpha$ represents the computational effectiveness of one information acquisition step. Suppose a bidder has a Gaussian valuation distribution $f_{\mu,\sigma}$ about her true value v for some good. Performing m consecutive steps of information acquisition as specified in Proposition 1 is equivalent to updating $f_{\mu,\sigma}$ to a new valuation distribution $f_{\mu',\sigma'}$, where (i) $f_{\mu',\sigma'}$ is also Gaussian, (ii) $\sigma' = \alpha^m\sigma$, and (iii) μ' is chosen according to a Gaussian density with mean μ and standard deviation $\sqrt{1-\alpha^{2m}}\sigma$.*

Proof. To prove this result, we need to show that the distribution of values from performing one step of information acquisition m times is the same as the distribution described above in Proposition 2. In other words, we want to show that

$$f_{\mu,\sigma}(x) = \int_{-\infty}^{\infty} f_{\mu,\sqrt{1-\alpha^{2m}}\sigma}(y) f_{y,\alpha^{2m}\sigma}(x) \, dy \tag{2.1.2}$$

holds for all m. This follows easily by induction on m, where (2.1.1) is used as the base case (see Appendix for proof). □

Equipped with a valuation model for bidders in which information acquisition is consistent, we can now characterize our representation of bidders in auctions. To start, we initialize all bidders at the same mean ν and standard deviation ρ. For an *uninformed* bidder i, we have her acquire one piece of information, giving her a new mean μ_0^i (chosen according to $f_{\nu,\sqrt{1-\alpha^2}\rho}$) and standard deviation $\sigma_0 = \alpha\rho$. So an uninformed bidder i enters an auction with mean μ_0^i and standard deviation σ_0. When considering a set of **uninformed bidders** we refer to their current means about their true value being distributed according to $f_{\nu,\sqrt{1-\alpha^2}\rho}$, which denote as γ.

On the other hand, at times we find it becomes necessary to consider bidders who are completely informed of their values. To keep our analysis consistent, we model an *informed* bidder j by having her theoretically acquire an infinite amount of information, i.e. $m \to \infty$. So an informed bidder j enters the auction with value v_j where v_j is chosen according to $f_{\nu,\rho}$ (since $\sqrt{1-\alpha^{2m}} \to 1$ as $m \to \infty$). Thus, when considering a set of **informed bidders**, we refer to their values being distributed according to $f_{\nu,\rho}$, which we now denote as ζ.

By representing bidders' beliefs about their true values with a Gaussian distribution , we have established a model for a bidder's valuation problem that can support multiple steps of information acquisition. From this point forward, we represent an *uninformed bidder i's current beliefs* about her valuation by the ordered pair (μ^i, σ^i), where μ^i is the mean and σ^i is the standard deviation of her Gaussian valuation distribution f.

2.1.2 Solving the Bidding Problem

A bidder faces the *bidding problem* when she must determine her optimal bid given information about her true value for the good, other bidders' bids, and the auction mechanism.

In this paper, we deal with auctions in which the dominant strategy for an informed, risk-neutral bidder is to bid "truthfully" according to her value v for the good. In particular, bidding "truthfully" means that (a) in a posted price auction, the bidder's dominant strategy is to accept a price p if and only if $v > p$ and to reject otherwise; (b) in a second-price sealed bid auction, the bidder's dominant strategy is to submit a bid $b^* = v$; and (c) in an ascending-price auction, the bidder's dominant strategy is to drop out once the price level p exceeds v. So, the optimal bid for an informed bidder in these auctions depends only on her own valuation, and not on the actions of other bidders. This simplifies the bidding problem because bidders need not model the bidding strategies of other bidders.

In this paper we deal with uninformed bidders that are uncertain about their private valuations. Our goal is to solve the bidding problem for these uninformed bidders.

2.1.3 Formulating the Information Acquisition Problem

The *information acquisition problem* for a bidder is to determine whether to acquire information in an auction and, if so, when and how much to acquire. On the one hand, acquiring information reduces a bidder's uncertainty about her true value for the good and so increases the accuracy of her bid; on the other hand, acquiring information is costly. Unlike the optimal bidding strategy, which depends only on a bidder's true value, the optimal information acquisition strategy may be much more complex and may depend on other bids, which in turn depend on whether other bidders can and do acquire information.

To formulate the problem of information acquisition, we use Russell and Wefald's general framework for belief refinement in bounded rational agents [16]. They present a fairly standard decision-theoretic, statistical approach and the mathematical foundations can be found in Berger [2]. Russell and Wefald consider an agent in an uncertain world, who takes both real-world "base-level" actions as well as "computational" actions to refine his beliefs. The value of a computational action is derived from its effect on the base-level action. Only if computation changes the base-level action to an action with higher utility does the computational action add value. If computation does not change the base-level action, even if it changes the agent's beliefs, it has no value.

Let S denote a computational action, and let S^m denote a sequence of computational actions. Also, let a denote the agent's current best base-level action based on its beliefs, and let a_{S^m} denote the agent's best base-level action after a sequence of S^m refinements, where $m \geq 1$. Define the *net value* of a computation action S to to be the change in utility of the best base-level action based on beliefs before and after the computation, minus the cost of computation:

$$V(S^m) = U(a_{S^m}) - U(a) - C(S^m)$$

where $U(a)$ is the utility of the state that follows base-level action a and $C(S^m)$ is the cost of a sequence of m computation actions.

In our analysis of information acquisition in auctions, we have the following correspondences:

- Agents correspond to bidders.
- Base-level actions correspond to bids.
- Computation actions correspond to information acquisition, and we denote a sequence of information acquisition steps as IA^m (as opposed to S^m).

- The state that follows a bid is determined by the auction mechanism and the other bids received. It also eventually defines the outcome of the auction (namely the bidder that wins the good and the price she pays).
- The utility of a state depends on whether the bidder wins the good, the price of the good, and the bidder's value for the good. The utility function corresponds to our risk-neutral bidders' von-Neumann Morgenstern utility for money, as introduced above. Bidder i has utility $U_i = v_i - p$ if it wins the good at price p, and 0 otherwise.
- The cost function of a single step of information acquisition is independent of the action taken, constant over time, and additive over multiple computations. So $C(IA) = c$ for some constant c and $C(IA^m) = mc$ where m is the number of information acquisition steps.
- The value of m steps of information acquisition (IA^m) then becomes

$$V(IA^m) = U(b^*_{IA^m}) - U(b^*) - mc \tag{2.1.3}$$

  where b^* is the bidder's current optimal bid before acquiring information and $b^*_{IA^m}$ is the bidder's optimal bid after m steps of information acquisition.

Notice that the information acquisition problem is difficult because of the uncertainty associated with the parameters that determine the optimal information acquisition strategy. There is uncertainty about the bidder's true value for the good; there is uncertainty about the outcome of information acquisition; and there is uncertainty about the state that follows a bid (because we don't know the bids of other bidders). So, a bidder cannot just apply (2.1.3) immediately to determine her optimal information acquisition strategy. She is only able to *estimate* the utility of her revised bid because it depends on the other bids received.

Because information acquisition necessarily involves uncertainty about its outcome, probability and decision theory are useful tools to apply when bidders can only estimate values. Following Russell and Wefald, let $\hat{Q}^S$ denote an *estimate* of the quantity Q obtained through computation S. We can then rewrite (2.1.3) as a new expression that bidders are able to compute. In particular, we determine an estimate of the value of m steps of information acquisition, $\hat{V}^{IA}(IA^m)$, as the estimated change in utility caused by the information acquisition minus the cost of acquiring information:

$$\hat{V}^{IA}(IA^m) = \hat{U}^{IA}(b^*_{IA^m}) - \hat{U}^{IA}(b^*) - mC. \tag{2.1.4}$$

A bidder can now use this equation to decide whether m steps of information acquisition is worthwhile by using information about the other bidders, the auction mechanism, and a model for her true private value for the good.

We use the above framework to formulate and solve the information acquisition problem for bidders in several auction mechanisms, which we describe in the next section.

2.2 Auction Formats

In this paper, we analyze and compare the optimal strategies for uninformed bidders in three different auction mechanisms, each with a single object for sale.

1. *Posted Price Sequential Auction*
   The auctioneer chooses a fixed price p and offers the price to each bidder in a random order as a take-it-or-leave-it offer. The good is sold to the first bidder (if any) that accepts the price. To achieve reasonable performance, we assume the seller is *knowledgeable* about the bidders in the market. Specifically, the seller chooses a competitive price p that maximizes revenue given bidders' hard valuation problems and his beliefs about the distribution of bidder values. When offered the good, bidders are given time to acquire information.

2. *Second-Price Sealed-Bid (Vickrey) Auction*
   Each bidder submits a sealed bid, and the item is sold to the highest bidder for the second-highest bid price. Bidders are given time to acquire information before submitting their bids.

3. *Asynchronous Ascending-Price (English) Auction*
   The auction maintains an ask price p, and bidders can always choose to bid at the current ask price. The auction also maintains a provisional allocation, indicating the current highest bidder, and increases p by some bid increment, $\epsilon > 0$, whenever a bid is received at the current price. The auction also maintains a list of "active" bidders who have not yet dropped out. All bidders are initially active. Whenever the auction enters a period of quiescence, with no bids received for a fixed period of time, the auctioneer picks, at random, one of the bidders not winning and tells her she must either submit a bid or leave the auction. When a bidder leaves, she is marked as not active and can never submit another bid. The auction terminates when there is only one active bidder remaining. This bidder wins the item at her final bid price. Bidders are able to acquire information at any time but are restricted by how much information they can acquire at any given time.

It is worth briefly explaining why we have chosen an asynchronous ascending-price auction over a synchronous one and why we have chosen to have the auctioneer intervene after a period of quiescence. We prefer the ascyhronous auction over the synchronous one because the former grants bidders a larger set of actions. In particular, while the synchronous auction requires each bidder to bid implicitly at every price level so long as she doesn't drop out, in the asynchronous auction a bidder can choose to remain active instead of either bidding or dropping out. This option to remain uncommitted allows the bidder to make better decisions about acquiring information. Furthermore, the motivation for having the auctioneer intervene is to restrict the need to introduce mixed-strategy equilibria into the bidders' optimal strategies. During a period of quiescence, each bidder prefers to remain active and wait for another bidder to take the risk, acquire information and bid the price up. Since all the bidders feel the same way, if the auctioneer does not interfere, the optimal strategy involves a mixed strategy between sitting tight and breaking the silence.

Following Parkes' terminology [11], we characterize the ascending-price and second-price sealed-bid auctions as *dynamic* since the price is determined after the bids are received, while the posted-price sequential auction is considered *static*. Furthermore, we describe the ascending-price and posted-price auctions as *progressive* since bidders receive feedback from the auction as they bid, while the second-price sealed-bid auction is *single-shot*. The posted-price is the only mechanism without dynamic price discovery and so the auctioneer needs to have knowledge about the distribution of bidders' values in the auction in order to achieve reasonable efficiency and revenue. Also, the ascending-price is the only mechanism that gives feedback about the actions of bidders to other bidders.

Given any auction mechanism, there are a few points of interest to consider and to determine whether they qualitatively affect our analysis. We list three of them here:

1. Does every bidder know the number of other bidders participating in the auction? Although this knowledge may not affect strategies in posted-price mechanisms, it can have a significant effect in any of the dynamic auctions.

2. In the ascending-price auction situation, do all bidders observe when another bidder drops out? This is referred to as overt versus secret dropout. Rezende [14] presents an analysis of the Japanese-style auction with secret dropout and claims that making dropout points visible does not qualitatively affect the analysis. Compte and Jehiel [4] on the other hand analyze the Japanese-style ascending-price auction where dropout points are overt and claim that this observability of dropout points is the single driving factor behind the preference of iterative auctions over sealed-bid auctions and thus that an ascending-price auction with secret dropout fails to achieve a better outcome than a sealed-bid auction. We come back to this question when we draw our conclusions about the different auctions.

3. When approaching his decision problem, does a bidder consider the bounded rationality of other bidders? Or does it assume that all the other bidders know their exact true valuations? We choose to focus on this point throughout the paper, as it is influential in determining optimal strategies in the dynamic auctions.

All these points consider the *knowledge* provided to bidders by the auction mechanism about the competition. As we will see, differences in the amount of knowledge provided is the main basis that generates different outcomes for different auction mechanisms.

2.3 Numerical Examples to Consider

It is helpful to describe realistic numerical examples to motivate the theory that follows in this paper. We outline a few here to keep in mind as the paper progresses. Recall the example described in the Introduction of the couple wanting to buy a new home. Having seen the home for the first time, the couple believes their value for the house is normally distributed with mean around \$10 (assume all figures in thousands), and standard deviation \$2.5. They also know that they can acquire additional information about the house for a cost. For example, they can consult an expert, buy some instructive guides to real estate, find out who has lived in the house previously, determine the quality of its construction, etc. Suppose that the couple knows that for each piece of information they acquire, they incur a fixed cost $c = \$0.01$ and reduce the standard deviation, i.e. their uncertainty about their true value, by the fraction $\alpha = 0.7$. So if they acquire m pieces of information, they incur a cost $\$0.01m$ and reduce their standard deviation by 0.7^m.

Now we will consider three scenarios, each of which involves the couple in a different auction mechanism:

1. *Posted Price.* The homeowner has heard the couple is interested in her home and offers them the opportunity to buy the home for $p = 8$. Although it is a take-it-or-leave-it offer, she kindly gives them a bit of time to make their decision. Should the couple use the time to acquire more information and gain certainty in their true value for the home? If so, how much information? After acquiring information, does the couple accept or reject the offer? We work this problem out in Chapter 3.

2. *Second-Price Sealed-Bid.* The home is being sold using a second-price sealed-bid auction format. Interested parties submit bids specifying the amounts they are willing to pay for the home, and the party with the highest bid pays a value equal to the second-highest bid. The couple has a few days to decide what bid to submit. Again,

how much more information should the couple acquire to gain certainty in their true value? Some? None? What bid should the couple ultimately submit? We work this problem out in Chapter 4.

3. *Ascending-Price Auction.* Now the home is being sold using an asynchronous ascending-price auction format, like that described above. The couple enters the auction as uninformed bidders with an initial distribution over their value for the home. Over the course of the auction, they have opportunities to take time and acquire pieces of information. When should they acquire information? How much information should they acquire? When should they bid? When should they drop out? We work through this problem in Chapter 5.

2.4 Outline

In the following three chapters, we analyze the equilibrium bidding and information acquisition strategies for uninformed bidders in each of the auction mechanisms. Although mechanism design cannot simplify the valuation problem, it can affect the efficiency of bidding and information acquisition.

We begin by looking at the posted-price auction and derive the optimal information acquisition strategy for an uninformed bidder. We then apply this result when looking at both second-price sealed-bid and ascending-price auctions. In our analyses, we first assume a basic model in which an uninformed bidder assumes all the other bidders are *informed* of their true values. We then extend our model to one in which the uninformed bidder assumes all the other bidders are *uninformed* as well. The advantage of the basic model is that it allows for an easy characterization of equilibrium behaviors. When a bidder considers that several buyers may also acquire information, however, the decision to acquire information becomes more complicated and is affected by others' decisions to acquire information as well.

Chapter 3

Posted-Price Sequential Auction

In this chapter, we formulate and solve the equilibrium bidding and information acquisition strategies for a bidder in a posted-price sequential auction. We then use this analysis of the posted-price auction as a foundation upon which to build our later analyses of dynamic auctions because the posted-price auction is an analytically simpler situation. Notice that we need only consider the scenario with one uninformed bidder because she is given a take-it-or-leave-it offer and so her decision problem does not depend at all on other bidders.

In a posted-price auction, when a bidder is offered the good, she is facing a fixed price and knows that if she accepts the price, she will receive the good for sure. There is no uncertainty about other bids or whether she will win. Thus, the posted-price auction isolates the uncertainty of information acquisition.

There are two problems to compute when a bidder i is offered the good: her optimal bidding strategy and her optimal information acquisition strategy. We assume that bidder i has a belief distribution (μ^i, σ^i) about her value v_i. In the analysis, we drop the i's for simplicity since we are dealing with a single bidder.

3.1 Optimal Bidding Strategy

A bidder that is no longer acquiring information has two options, either to accept the price or to reject the price. Given a fixed price p, her expected utility for accepting the price is

$$\begin{aligned} \hat{U}(accept) &= E[v] - p \\ &= \mu - p, \end{aligned}$$

while her expected utility for rejecting the price is

$$\hat{U}(reject) = 0. \tag{3.1.1}$$

Therefore, the bidder's optimal bidding strategy is to accept whenever $\hat{U}(accept) \geq \hat{U}(reject)$ or $\mu - p \geq 0$, and to reject otherwise. The expected utility of this optimal bidding strategy is $\hat{U}(b^*((\mu, \sigma), p)) = \max[0, \mu - p]$.

3.2 Optimal Information Acquisition Strategy

When faced with a fixed price p, a bidder may choose to exercise his ability to acquire information. Recall that there is a fixed cost c per step of information acquisition. We begin by considering the case where $\mu \geq p$.

Let m consecutive steps of information acquisition be represented by IA^m. Before acquiring information, since $\mu \geq p$, the best action is to accept. The expected utility of accepting is always $\mu - p$. After acquiring information, let the new mean be μ_m. If $\mu_m \geq p$, the revised best action is to accept. If $\mu_m < p$, the revised best action is to reject. Then the expected utility of the revised best action after information acquisition is $\mu_m - p$ when $\mu_m \geq p$ and 0 when $\mu_m < p$. Following the intuition of Russell and Wefald, we expect the information acquisition to have value only if it changes the best action. In our case, the best action changes when $\mu_m < p$.

Let the bidder's expected gain from IA^m when the price is p (remember that $\mu \geq p$) be denoted by the function $H^{\mu \geq p}(m, p)$, where $m > 0$. Mathematically, using 2.1.4, $H^{\mu \geq p}(m, p)$ is the integral over the possible new mean μ_m of the difference between the utility of the revised best action *after* acquiring information and the utility of the best action *before* acquiring information given that the new mean is μ_m, all minus the cost mc of acquiring information:

$$\begin{aligned} H^{\mu \geq p}(m, p) &= \left(\int_{-\infty}^{p} f_{\mu, \sqrt{1-\alpha^{2m}}\sigma}(\mu_m)(0 - (\mu_m - p))\, d\mu_m \right) \\ &+ \left(\int_{p}^{\infty} f_{\mu, \sqrt{1-\alpha^{2m}}\sigma}(\mu_m)((\mu_m - p) - (\mu_m - p))\, d\mu_m \right) - mc \\ &= \left(\int_{-\infty}^{p} f_{\mu, \sqrt{1-\alpha^{2m}}\sigma}(\mu_m)(p - \mu_m)\, d\mu_m \right) - mc. \end{aligned} \tag{3.2.1}$$

We know that the Gaussian density f of mean μ and standard deviation σ is given by

$$f_{\mu,\sigma}(x) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\mu)^2/2\sigma^2}.$$

To simplify our notation, we introduce $s(m,\sigma) = (1-\alpha^{2m})\sigma^2$. Plugging those last two equations into 3.2.1, we can simplify our expression to the following:

$$\begin{aligned} H^{\mu \geq p}(m,p) &= \left(\int_{-\infty}^{p} (p-\mu_m) \frac{1}{\sqrt{2\pi s(m,\sigma)}} e^{\frac{-(\mu_m-\mu)^2}{2s(m,\sigma)}} \, d\mu_m \right) - mc \\ &= \frac{1}{\sqrt{2\pi s(m,\sigma)}} \int_{-\infty}^{p} (p-\mu_m) e^{\frac{-(\mu_m-\mu)^2}{2s(m,\sigma)}} \, d\mu_m - mc \\ &= \frac{s(m,\sigma) e^{\frac{-(p-\mu)^2}{2s(m,\sigma)}} + \frac{1}{2}\sqrt{2\pi s(m,\sigma)}(p-\mu)(1+\Psi(\frac{p-\mu}{\sqrt{2s(m,\sigma)}}))}{\sqrt{2\pi s(m,\sigma)}} - mc \end{aligned}$$

where $\Psi(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} \, dt = 2\Phi(\sqrt{2}x) - 1$ and where $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} e^{-t^2/2} \, dt$ is the standard normal cumulative distribution function, whose values have been widely documented in tables.

Now consider the case where $\mu \leq p$. Proceeding in an identical manner, we obtain a similar result:

$$H^{\mu \leq p}(m,p) = \frac{s(m,\sigma) e^{\frac{-(\mu-p)^2}{2s(m,\sigma)}} + \frac{1}{2}\sqrt{2\pi s(m,\sigma)}(\mu-p)(1+\Psi(\frac{\mu-p}{\sqrt{2s(m,\sigma)}}))}{\sqrt{2\pi s(m,\sigma)}} - mc$$

And hence for all p, we have the following result:

$$H(m,p) = \frac{s(m,\sigma) e^{\frac{-(p-\mu)^2}{2s(m,\sigma)}} - \frac{1}{2}\sqrt{2\pi s(m,\sigma)}|p-\mu|(1-\Psi(\frac{|p-\mu|}{\sqrt{2s(m,\sigma)}}))}{\sqrt{2\pi s(m,\sigma)}} - mc. \tag{3.2.2}$$

Notice that $H(m,p)$ is symmetric around μ given a fixed m.

Consequently, a necessary condition for a bidder to acquire information m times when faced with price p is if $H(m,p) > 0$. However, the strategy of acquiring information m consecutive times is dominated by the strategy of acquiring information once and then using that information in deciding whether to acquire additional information, because it is possible that the new information makes future information acquisition unprofitable. Thus, we have the following proposition.

Proposition 3. *In a posted-price sequential auction, when a bidder is offered a good at price p, the optimal information acquisition strategy is to acquire information as long as there is some $m \geq 1$ number of information acquisition steps such that $H(m,p) > 0$. After each step of information acquisition, however, the bidder re-confirms that there still exists $m \geq 1$ such that $H(m,p) > 0$. Otherwise, she stops acquiring information.*

With that result in mind, it is natural to consider whether there exists an optimal $m^*(p) = m^*$ that maximizes $H(m,p)$ for a given p. If there is such an m, then the bidder need only check whether $H(m^*,p) > 0$ in determining whether to acquire information. To answer such questions, we look at the function $H(m,p)$.

3.2.1 What is the shape of $H(m,p)$ for a fixed p ?

To begin, we look at the shape of $H(m,p)$ for a fixed p, which we denote as $H(m)$. Recall that $H(m)$ is defined for $m > 0$. In particular, we want to answer the following questions:

- Is there a scenario in which $H(m)$ is always decreasing (in which case we need only check if IA^1 is worthwhile? Or always increasing (in which case we may not ever stop acquiring information)?
- Does $H(m)$ have a single global maximum? In other words, is there an optimal m^*?

To answer those questions, consider the derivative of $H(m)$ with respect to m. Taking derivatives, we find that

$$\frac{\partial H(m)}{\partial m} = \frac{\sigma\alpha^{2m}\log(\frac{1}{\alpha})e^{\frac{-(p-\mu)^2}{2(1-\alpha^{2m})\sigma^2}}}{\sqrt{2\pi(1-\alpha^{2m})}} - c. \tag{3.2.3}$$

Let the first term be represented by $\chi(m,p) = \chi(m)$. We know that (a) if $\chi(m) \leq c$ for all m, then $H(m)$ is monotonically decreasing; (b) if $\chi(m) \geq c$ for all m, then $H(m)$ is monotonically increasing; (c) otherwise, $H(m)$ is increasing for some values m and decreasing for other values m. Since $\alpha < 1$ and $\sigma > 0$, we note that $\chi(m)$ is always positive (for $m > 0$).

Before we continue with the analysis, we first single out the case when $\mu = p$. In this case, we can see that $\chi(m)$ explodes ($\rightarrow \infty$) as $m \rightarrow 0$, but strictly decreases asympotically to 0 as $m \rightarrow \infty$ (see Appendix for proof). Thus, exactly one solution m^* exists such that $\chi(m^*) = c$ and hence $\frac{\partial H(m)}{\partial m} = 0$. Since $\chi(m)$ is strictly decreasing, we know that m^* is in fact the global maximum of $H(m)$ when $\mu = p$. In fact, we can solve the equation $\chi(m^*) = c$ for m when $\mu = p$, obtaining

$$m^*_\mu = \frac{1}{2}\log_\alpha\left[\frac{c\pi}{\sigma^2(\log\alpha)^2}\left(\sqrt{c^2 + \frac{2\sigma^2(\log\alpha)^2}{\pi}} - c\right)\right]. \tag{3.2.4}$$

Since a bidder must acquire information at least $m = 1$ times, a bidder's optimal information acquisition strategy is the following: *when $\mu = p$, if $m^* \geq 1$, then acquire information only if $H(m^*,p) > 0$; else, acquire information only if $H(1,p) > 0$.*

Now we consider the case where $\mu \neq p$. When we look at the limits, although $\chi(m)$ is not defined at 0, as $m \to 0$, $\chi(m) \to 0$ since the exponential term dominates. Also, as $m \to \infty$, $\chi(m) \to 0$ since $\alpha^{2m} \to 0$. Since $\chi(m)$ does not explode, it is possible to satisfy (a) and have $H(m)$ strictly decreasing. However, since $\chi(m)$ is continuous, there is always a m for which $\chi(m) < c$ and hence (b) is impossible and $H(m)$ is never strictly increasing.

Consider the function $\chi(m)$. Does it have a global maximum m^{**}? If so, we need only compare $\chi(m^{**})$ to c to tell if $H(m)$ is strictly decreasing. We take the derivative of $\chi(m)$ and obtain

$$\frac{\partial \chi(m)}{\partial m} = \frac{-\sigma^2 \alpha^{2m} \log(\alpha)^2 e^{\frac{-(p-\mu)^2}{2(1-\alpha^{2m})\sigma^2}}}{\sigma\sqrt{2\pi(1-\alpha^{2m})}} \left(\frac{\alpha^{2m}}{1-\alpha^{2m}} - \frac{\alpha^{2m}(p-\mu)^2}{(1-\alpha^{2m})^2\sigma^2} + 2 \right).$$

Setting that expression equal to 0, we obtain

$$\frac{\alpha^{2m^{**}}\sigma^2}{(1-\alpha^{2m^{**}})\sigma^2} + 2 = \frac{\alpha^{2m^{**}}(p-\mu)^2}{(1-\alpha^{2m^{**}})^2\sigma^2}$$
$$\Longrightarrow \alpha^{4m} - \alpha^{2m}(3 + (\frac{p-\mu}{\sigma})^2) + 2 = 0. \tag{3.2.5}$$

Using the quadratic formula as well as the fact that $\alpha^{2m} < 1$ for all m (see Appendix for proof), we conclude that if $\mu \neq p$,

$$m^{**} = \frac{1}{2}\log_\alpha \left[\frac{1}{2} \left(3 + (\frac{p-\mu}{\sigma})^2 - \sqrt{\left(3 + (\frac{p-\mu}{\sigma})^2\right)^2 - 8} \right) \right]. \tag{3.2.6}$$

Since $\chi(m)$ is always positive, $\chi(m) \to 0$ both as $m \to 0$ and as $m \to \infty$, and $\frac{\partial \chi(m)}{\partial m} = 0$ has a unique solution, $\chi(m^{**})$ is the unique global maximum of $\chi(m)$ for $m > 0$. So we have the following result.

Lemma 4. *Assume $\mu \neq p$. If $\chi(m^{**}) \leq c$, then $H(m)$ is monotonically decreasing in m. Otherwise, there are exactly two solutions for m in $\frac{\partial H(m,p)}{\partial m} = 0$, namely $m_1^* \in (0, m^{**})$ and $m_2^* \in (m^{**}, \infty)$, where m_1^* defines the minimum of $H(m)$ and m_2^* defines the maximum.*

Proof. Since the global maximum of $\chi(m)$ is defined at m^{**}, if $\chi(m^{**}) \leq c$, then $\chi(m) \leq c$ for all m and thus $H(m)$ is always decreasing in m.

If $\chi(m^{**}) > c$, since $\chi(m)$ is continuous on $(0, \infty)$, approaches 0 both as $m \to 0$ and as $m \to \infty$, *and* has only one extremum, it follows from the Intermediate Value Theorem that there must be one $m_1^* \in (0, m^{**})$ such that $\chi(m_1^*) = c$ and one $m_2^* \in (m^{**}, \infty)$ such that $\chi(m_2^*) = c$. Since the condition that $\chi(m) = c$ is equivalent to the condition that $\frac{\partial H(m)}{\partial m} = 0$, there are exactly two solutions for either equation.

But now we have all the facts we need to determine the shape of $H(m)$, assuming it is not monotone! In particular, on the interval $(0, m_1^*)$, we know $\frac{\partial H(m)}{\partial m} < 0$ and so $H(m)$ is decreasing. In the interval (m_1, m_2), $\frac{\partial H(m)}{\partial m} > 0$ and so $H(m)$ is increasing. Finally, in the interval (m_2, ∞), $\frac{\partial H(m)}{\partial m} < 0$ and so $H(m)$ is decreasing again. Thus we conclude that m_1^* defines the minimum of $H(m,p)$ and m_2^* defines the maximum. Both m_1^* and m_2^* can easily be solved for numerically, although they lack explicit analytic forms. □

If $H(m)$ is not monotone, the maximum value of $H(m)$ is achieved at $H(m_2)$ and so $m^* = m_2^*$. Thus, since a bidder must acquire at least $m = 1$ steps of information, the bidder's optimal information acquisition strategy is the following: *when $\mu \neq p$, if $m^* \geq 1$, then acquire information if and only if $H(m^*) > 0$; else, acquire information if and only if $H(1) > 0$.* Notice that it may be sufficient to check if $H(m^{**}) > 0$ since if that condition holds then it follows that $H(m^*) > 0$.

We have fully characterized the shape of $H(m,p)$ for a fixed price p, given $m > 0$. Figures **??** illustrate the four different shapes $H(m)$ can take.

Now We are able to fully define m^*. In particular,

$$m^* = \begin{cases} m_\mu^* & \mu = p \\ m_2^* & \mu \neq p \end{cases} \tag{3.2.7}$$

Now, we obtain a complete characterization of a bidder's optimal information acquisition strategy in a posted-price auction.

Proposition 5. *Let a bidder be offered a good for price p in a posted-price sequential auction. She acquires information if and only if (i) $m^* \geq 1$ and $H(m^*, p) > 0$ or (ii) $m^* < 1$ and $H(1,p) > 0$. After each step of information acquisition, the bidder re-confirms that information acquisition is still profitable, i.e. that her beliefs still satisfy either (i) or (ii). Otherwise, she stops acquiring information.*

3.2.2 What else do we know about $H(m,p)$?

For the sake of future discussion, we also consider the shape of $H(m,p)$ for a fixed value of m, which we denote as $H(p)$. Taking the derivative of $H(p)$ with respect to p, we obtain

$$\frac{\partial H(p)}{\partial p} = I_0(\mu,p)\frac{1}{2}\left(1 - \Psi\left(\frac{|p-\mu|}{\sqrt{2s(m,\sigma)}}\right)\right)$$

where $I_0(\mu,p) = 1$ if $\mu \geq p$ and $I_0(\mu,p) = -1$ if $\mu \leq p$. Thus, given μ, if $\mu \geq p$, $H(p)$ is increasing and if $\mu \leq p$, $H(p)$ is decreasing. Furthermore, $H(p)$ is perfectly symmetric around μ, since for $p < \mu$ $H(p) = H(2\mu - p)$. See Figure.

Hacquire.eps

Figure 3.1: $\mu \neq p$, $\chi(m^{**}) > c$, bidder acquires

Hnoacquire.eps

Figure 3.2: $\mu \neq p$, $\chi(m^{**}) > c$, bidder does not acquire

Hchilessc.eps

Figure 3.3: $\mu \neq p$, $\chi(m^{**}) \leq c$, bidder acquires

Hmuequalsp.eps

Figure 3.4: $\mu = p$, bidder acquires

Hofp.eps

Figure 3.5: H(p) with $\mu = 10$

Result 6. *For fixed μ and m, $H(p)$ is symmetric around μ.*

Next, we look at the behavior of $m^*(p)$ for a fixed μ as $|p - \mu|$ varies. Consider the case where $p < \mu$. Define p_1 and p_2 such that $p_1 < p_2 < \mu$. Let $m_1^* = m^*(p_1)$ be the m that maximizes $H(m, p_1)$ and let $m_2^* = m^*(p_2)$ be the m that maximizes $H(m, p_2)$. We argue that it must be that $m_1^* < m_2^*$.

From our previous discussion about the shape of $H(m, p)$, we know that $\chi(m_2^*, p_2) = c$ and $\chi(m_1^*, p_1) = c$. Also, by looking at our expression for $\chi(m, p)$, we see that for a fixed m, if $p < \mu$, as p increases $\chi(m, p)$ increases as well. Thus, since $p_2 > p_1$, we find $\chi(m_1^*, p_2) > c$. But recall the shape of $\chi(m)$! It reaches its maximum at m^{**} $(< m^*)$ and then is monotonically decreasing for $m > m^{**}$. So if $m > m_2^*$, then $\chi(m, p_2) < \chi(m_2^*, p_2) = c$. But since $\chi(m_1^*, p_2) > c$, we conclude that it must be the case that $m_2^* > m_1^*$. The case where $\mu > p$ is identical and we have the following result.

Result 7. *For fixed μ, as p moves further away from μ, i.e. $|p - \mu|$ increases, the optimal number of information acquisition steps m^* decreases. In other words, m^* and $|p - \mu|$ are inversely related.*

Finally, it is also interesting to look at the behavior of the actual value of $H(m^*, p)$ for a fixed μ as $|p - \mu|$ varies. Although on first glance it may seem that the Envelope Theorem is required, we make use of our previous results to come to the conclusion in a much simpler fashion.

Again, we assume without loss of generality that $p < \mu$ (the other case is identical). Recall p_1, p_2, m_1^* and m_2^* as defined above, where $p_1 < p_2$ and $m_1^* < m_2^*$. Since (a) $H(p)$ is decreasing in $|p - \mu|$ for a fixed m and (b) $H(m^*, p) > H(m, p)$ for any $m \neq m^*$, we have the following inequalities:

$$
\begin{aligned}
H(m_2^*, p_2) &> H(m_1^*, p_2) \\
H(m_1^*, p_2) &> H(m_1^*, p_1),
\end{aligned}
$$

which in turn imply that

$$
H(m_2^*, p_2) > H(m_1^*, p_1).
$$

Thus we conclude that as $|p - \mu|$ increases, $H(m^*, p)$ decreases.

Also, looking at the expression for $\chi(m, p)$, since $\chi(m, p) = \chi(m, p')$ if $|p' - \mu| = |p - \mu|$, we find that $H(m^*, p)$ is symmetric around μ.

Result 8. *For fixed μ and p, as p moves further away from μ, i.e. $|p-\mu|$ increases, the expected gain from acquiring the optimal number of steps of information, $H(m^*,p)$, decreases. In other words, $H(m^*,p)$ and $|p-\mu|$ are inversely related.*

Taking those results together, we discover that for a fixed μ, $H(m^*,p)$ reaches its maximum value at $H(m^*_\mu,\mu)$ and that the m^*_μ such that $m^*_\mu = \arg\max_m H(m,\mu)$ is the largest possible number of information acquisition steps a bidder will take. So, given any bidder and a posted price p, we know that her optimal number of information acquisition steps is between m^{**} and m^*_μ. So,

Result 9. *For a fixed p, the maximum number of optimal information acquisition steps for any bidder will be m^*_μ, as detailed in (**??**).*

3.3 Numerical Example Revisited

We return to the situation outlined in Chapter 2 about the couple who is being offered the home for $p = 8$. Recall that the couple has beliefs (μ_0, σ_0) about their value, where $\mu_0 = 10$ and $\sigma_0 = 2.5$. Also, recall that they can acquire m pieces of information for a cost of $0.01m$ and they know that their computational effectiveness, $1-\alpha$, is equal to 0.3 (i.e. $\alpha = 0.7$). Should the couple acquire information? Should they accept the price? We can now answer these questions using the equilibrium strategies outlined above.

First, notice that $\mu \neq p$. So, we plug our values into the expression for m^{**} and find out within what range we are looking for m^*. We find that

$$m^{**} = 0.5522.$$

For interest's sake, we also calculate $m^*_\mu = 5.026$. Thus, the optimal m^* number of deliberations lies in the interval $(0.5522, 5.026)$. Recall that we want to find $m^* \in (0.5522, 5.026)$ such that $\chi(m^*) = c$, i.e.

$$\frac{2.5 * 0.7^{2m^*} \log(1/0.7) e^{\frac{-0.32}{1-0.7^{2m}}}}{\sqrt{2\pi(1-0.7^{2m})}} = c.$$

Evaluating numerically, we obtain that

$$m^* = 4.5125.$$

Since $H(4.5125, 8) = 0.241 > 0$, the couple decides to acquire information once. They then re-evaluate whether to acquire information again. Using a simulator, we obtain that

$\mu_1 = 11.51$ and we know $\sigma_1 = \alpha\sigma_0 = 1.75$. Since $\mu \neq p$, we follow the same process and determine that $m^* = 3.571$. Since $H(3.571, 8) = -0.0245 < 0$, the couple decides to stop acquiring information.

Since $\mu_1 = 11.51 > p = 8$, the couple's optimal bidding strategy is to accept the offer. Their total gain from the purchase is $\mu_1 - p - 0.01 = 3.5$.

3.4 Looking Ahead

In any situation where only one bidder is acquiring information and must choose an action, she makes her decision to acquire information based on what she expects the "posted price" she is facing to be. Once this posted price is estimated, her optimal information acquisition strategy is to follow the strategy outlined above for a posted-price sequential auction. We use this idea when looking at more complex auctions in the following two chapters.

Chapter 4

Second-Price Sealed-Bid Auction

In this chapter, we formulate and solve the equilibrium bidding and information acquisition problems for a bidder in a second-price sealed-bid (SPSB) auction. Our analysis of the posted-price sequential auction in the previous chapter comes in handy as a useful tool.

In a SPSB auction, a bidder must decide what to bid. On first glance, it may not seem that she is facing any sort of fixed price. Instead, she wins the good only if her bid is greater than the bids of all the other bidders, and pays a price equal to the highest bid other than her own. In effect, however, she does face a "posted price" equal to the value of the highest of the other bids. This value depends on both the bidding strategies and potential information acquisition strategies of the other bidders.

In this chapter, we first formulate a bidder's optimal bidding strategy. This is straightforward since a bidder's optimal bidding strategy is independent of her beliefs about the other bidders' strategies. We then look at a bidder's optimal information acquisition strategy, considering two cases: (a) when a bidder believes that all the other bidders are *informed* of their exact values and hence do not acquire information, and (b) when a bidder believes that all the other bidders are *uninformed* and can also acquire information. That distinction is important because a bidder's decision to acquire information is affected by others' bids and others' decisions to acquire information, which in turn are affected by how informed the other bidders are.

4.1 Optimal Bidding Strategy

The bidder must decide the optimal bid b^* to submit in this auction. We know the SPSB auction is incentive compatible for informed bidders because the value of a bid b does not affect the price that a bidder pays for the good, but instead indicates the maximum price the bidder will accept, so long as that price is bid by another bidder. In fact, the SPSB auction is strategyproof [18]: an informed bidder's optimal bidding strategy is to bid her true value $b^* = v$.

The incentive-compatibility of the auction extends to a system where bidders' values are uncertain: the optimal bidding strategy is for a bidder to bid her expected value of the good (see proof in Appendix). So, the optimal bid for a bidder that is no longer acquiring information is $b^*(\mu, \sigma) = \mu$. Notice that this bid is optimal *regardless of the distribution over outside bids.* Since this optimal bid always has strictly positive expected value (she never pays more than her mean), a bidder should always bid.

4.2 Optimal Information Acquisition Strategy

In a SPSB auction, the value of acquiring information is computed as the increase in utility from placing a different bid minus the cost of acquiring information. In a sealed-bid auction, the value of acquiring information is uncertain because it depends on (a) whether the bidder will win the good, and (b) the price that the bidder will pay for the good. Both depend on the others' bidding and information acquisition strategies.

4.2.1 One Uninformed Bidder

Consider an auction with $n + 1$ bidders. Without loss of generality, let bidder 0 be the uninformed bidder. She assumes that all other bidders know their true values for the good and hence will not acquire information. Thus she only needs to consider their bidding strategies.

We call these bidders 1 through n the "informed bidders" and denote bidder i's value by v_i. For simplicity's sake, we index the bidders such that $v_1 > v_2 > \ldots > v_n$. According to the optimal bidding strategy for informed bidders in a SPSB auction, their bids will be equal to their values. Letting θ_i denote the bid of bidder i, we have $\theta_i = v_i$ for all $i \in \{1, 2, .., n\}$.

Assume that bidder 0 has a belief about the *distribution* from which the informed bidders' valuations are drawn. So we say she is *knowledgeable.* In particular, according to our model of informed bidders, she assumes that the density of that distribution is ζ with full support on $[\underline{\theta}, \bar{\theta}]$, where $\underline{\theta} \geq 0$. Recall from Chapter 2 that ζ is Gaussian with mean ν and standard deviation ρ.

If the informed bidders follow their optimal bidding strategies, bidder n is essentially facing a "posted price" of $p = \theta_1$, i.e. the highest order statistic of the n informed bidders' values. Let g denote the distribution of θ_1 (or equivalently p). From statistics, we know that the distribution of the highest order statistic of n values drawn from a distribution with density f and cumulative distribution function F is

$$f^{(n)}(x) = nf(x)F^{n-1}(x). \tag{4.2.1}$$

We can then write the distribution of p to be

$$g(p) = n\zeta(p)\left(\int_{-\infty}^{p} \zeta(t)\,dt\right)^{n-1}. \tag{4.2.2}$$

Since $\zeta(t)$ is Gaussian, we can represent the integral in terms of the standard normal cumulative distribution function $\Phi(z)$, because for any Gaussian cumulative distribution function F with mean μ and standard deviation σ, $F(x) = \Phi(\frac{x-\mu}{\sigma})$. So we obtain

$$g(p) = \frac{n}{\sqrt{2\pi}\rho} e^{\frac{-(p-\nu)^2}{2\rho^2}} \Phi\left(\frac{p-\nu}{\rho}\right)^{n-1}. \tag{4.2.3}$$

Now we can write down a formula for the value of acquiring information. The following is the expression for the expected utility of a bid b, given $g(p)$ and a bidder with belief distribution f over her value:

$$\hat{U}(b) = \int_{p=0}^{b} g(p) \int_{-\infty}^{\infty} (v-p)f(v)\,dv\,dp. \tag{4.2.4}$$

Let m consecutive steps of information acquisition be represented by IA^m. Let the bidder have beliefs (μ, σ) about her value before acquiring information, and let her have beliefs $(\mu_m, \alpha^m\sigma)$ after IA^m. So, according to the optimal bidding strategy, before IA^m, her best action is to bid $b^* = \mu$, and after IA^m, her revised best action is to bid $b^*_{IA^m} = \mu_m$.

Denote the bidder's expected gain from IA^m as the function $J_p(m)$, and recall that each step of information acquisition costs c. Mathematically, using 2.1.4, $J_p(m)$ is the difference between the utility of b^* and the utility of $b^*_{IA^m}$ given that the new mean is μ_m, all minus the cost mc of acquiring information:

$$
\begin{aligned}
J_p(m) &= \left(\int_{p=0}^{\mu} g(p) \int_{-\infty}^{p} f_{\mu,\sqrt{(1-\alpha^{2m})}\sigma}(\mu_m)(0-(\mu_m - p))\, d\mu_m\, dp\right) \\
&\quad + \left(\int_{p=\mu}^{\infty} g(p) \int_{p}^{\infty} f_{\mu,\sqrt{(1-\alpha^{2m})}\sigma}(\mu_m)((\mu_m - p) - 0)\, d\mu_m\, dp\right) \\
&\quad -mc \qquad (4.2.5) \\
&= \int_{\mu}^{\infty} g(p) H^{\mu \le p}(m,p)\, dp + \int_{0}^{\mu} g(p) H^{\mu \ge p}(m,p)\, dp \\
&= \int_{0}^{\infty} g(p) H(m,p)\, dp \qquad (4.2.6) \\
&= E_p[H(m,p)]. \qquad (4.2.7)
\end{aligned}
$$

Obviously, a necessary condition for a bidder to acquire information m times in the SPSB auction is if $J_p(m) > 0$. However, as we established in the previous chapter when discussing posted price auctions, the strategy of acquiring information m consecutive times is dominated by the strategy of acquiring information once and then using that information to decide whether to acquire additional information. Thus, we have the following proposition.

Proposition 10. *In a second-price sealed-bid auction, when a bidder believes he is facing informed bidders whose maximum value, p, is distributed according to density $g(p)$, the optimal information acquisition strategy is to acquire information as long as there is some $m \ge 1$ number of information acquisition steps such that $J_p(m) > 0$. After each step of information acquisition, the bidder re-confirms that there exists $m \ge 1$ such that $J_p(m) > 0$. Otherwise, she stops acquiring information.*

With that result in mind, it is natural to consider whether there exists an optimal m_J^* such that

$$m_J^* = \arg\max_m J_p(m)$$

for a given distribution $g(p)$. If there is such an m_J^*, then the bidder need only check whether $J_p(m_J^*) > 0$ in determining her optimal acquisition strategy. To answer such questions, we look at the function $J_p(m)$.

First, we note that $J_p(m)$ is the expected value of $H(m,p)$ over p given $g(p)$. Intuitively we want to find an m_J^* such that “on average” if p is distributed according to $g(p)$, $H(m_J^*, p)$ is as large as possible. Unfortunately, because of the complexities of both $g(p)$ and $H(m,p)$, although $J_p(m)$ is easily evaluated numerically for any given m and $g(p)$, it is difficult to deal with analytically (e.g. we cannot take a derivative with respect to m and set it equal to 0). As a result, we propose a conservative approach for estimating whether there exists an $m \ge 1$ such that $J_p(m) > 0$.

Approach to Estimate Optimal Information Acquisition Strategy

Our approach is first to calculate the expected value of p given $g(p)$, i.e.

$$E[p] = \int_0^\infty pg(p)\,dp$$

We then proceed as we would in the posted price auction, using $E[p]$ as the posted price and obtaining the $m^* = m^*(E[p])$ such that $H(m, E[p])$ is maximized. Next we evaluate $J_p(m^*) > 0$ and determine whether we should acquire information. Notice that $J_p(m^*) > 0$ is a *sufficient* condition for acquiring information (because then we know there exists *some* $m \geq 1$ such that $J_p(m) > 0$), but it is not necessary. Thus the danger is that we may fail to acquire information when we should.

So now, how good of an estimate is m^*? Essentially, we need to look at the two functions $E[H(m,p)]$ and $H(m, E[p])$, since

$$\begin{aligned} m_J^* &= \arg\max_{m\geq 1} E[H(m,p)] \text{ and} \\ m^* &= \arg\max_{m\geq 1} H(m, E[p]). \end{aligned}$$

We know that $E[H(m,p)] = H(m, E[p])$ when both $g(p)$ and $H(m,p)$ are symmetric around the same mean μ. Unfortunately, $g(p)$ is not symmetric. Instead, it is skewed to the left.[1] So, informally the "hump" of $g(p)$ is to the left of $E[p]$.

Recall that both the optimal $m^*(p)$ for a given p as well as the value of $H(m^*, p)$ increases the closer p gets to μ. Thus, we deduce generally that when $E[p] - \mu > 0$ our approach under-compensates for the high values of $H(m,p)$ achieved when p is near μ and m^* is typically an *underestimate* of m_J^*. Meanwhile, when $E[p] - \mu < 0$, our approach over-compensates for the high values of $H(m,p)$ achieved when p is near μ and so m^* is typically an *overestimate* of m_J^*.

Nevertheless, given the information we have, intuitively the above is a very reasonable approach to approximating m^* and we expect that $J(m^*)$ will in most instances have the same sign as $J(m_J^*)$. Notice that in this setup the uninformed bidder will only under-acquire information. This happens when $|m^* - m_J^*|$ is large enough such that $J_p(m^*) < 0$ but $J_p(m_J^*) > 0$. In this case it may be practical to search values between $m^* - \delta$ and $m^* + \delta$ for some integer δ determined by the expected error of m^*, and see whether there exists some m in that range such that $J_p(m) > 0$. Doing so will diminish the probability of bidders failing to acquire information when they should.

[1] This can be easily verified by graphing $g(p)$ for any value of n, ν, and ρ

Proposition 11. *In a second-price sealed-bid auction, when a bidder believes he is facing informed bidders whose maximum value p is distributed according to density $g(p)$, a bidder's optimal information acquisition strategy can be approximated by first calculating $E[p]$ and determining m^* using the algorithm outlined in 5. If $J_p(m^*) > 0$, then the bidder acquires information. After each step of information acquisition, the bidder re-confirms that $J_p(m^*) > 0$ still holds. Otherwise, she stops acquiring information. Note that this is a* conservative *approach in the sense that using this approximate strategy will never cause the bidder to acquire too much information.*

Remark 4.2.1. Notice that the bidders' optimal information acquisition strategy depends entirely on her having beliefs about the distribution of p, i.e. $g(p)$ (which requires the knowledge of both ζ and n, the number of bidders). Other than that outside knowledge, she is given no other guidance by the SPSB auction as to whether she should acquire information. Thus, an uninformed bidder that does not know $g(p)$ will have to follow some sub-optimal best-case or ad-hoc strategy to acquire information.

4.2.2 All Uninformed Bidders

In the SPSB auction, a bidder does not get any dynamic feedback on the values of the other bidders. What happens when she believes that *all* bidders are now uninformed?

In this situation, according to our model laid out in Chapter 2, the means of the current belief distributions for the uninformed bidders are distributed according to a γ, a Gaussian with mean ν and standard deviation $\sqrt{1-\alpha^2}\rho$.

Again let us assume we have $n+1$ bidders in the auction, each of which we index by $i \in \{0, 1, \ldots, n\}$. Each bidder is uninformed, can acquire information, and believes that all the other bidders are also uninformed and can acquire information. Assume that each bidder i enters the auction with an initial valuation distribution (μ_0^i, σ_0), where the μ_0^i are drawn according γ, and where $\sigma_0 = \alpha\rho$ is the same across the bidders.

So, the question we consider is the following: for any given bidder i, how does her information acquisition strategy change once she knows other bidders are also uninformed?

When all the bidders are completely uninformed, their values are distributed according to γ. When all the bidders are informed, their values are distributed according to $f_{\nu,\rho} = \zeta$. What happens when bidders acquire information? For the sake of argument, we assume a symmetric equilibrium in which every bidder acquires the same amount (m) of information. The consistency of our model implies that the bidders' values are now distributed according to $f_{\nu,\sqrt{1-\alpha^{2m}}\rho}$. Thus, as bidders acquire information symmetrically, the distribution over their values maintains the same mean μ but the standard deviation increases. Let the expected maximum value of the n bids be denoted by θ_1. Hence, we have the following proposition.

Proposition 12. *In a symmetric equilibrium where every bidder begins as being uninformed and acquires the same amount of information, the expected maximum value of the n bids, $E[\theta_1]$, increases as bidders acquire information. When bidders are fully informed, $E[\theta_1]$ is largest.*

Proof. Recall that the formula for the distribution of the highest order statistic of n values drawn from a distribution with density f and cumulative distribution function F is

$$f^{(n)}(x) = nf(x)F^{n-1}(x).$$

We also note the following fact from statistics.

Fact 13. *Let f and g be Gaussian distributions where both have mean μ but f has standard deviation σ and g has standard deviation ρ, where $\sigma > \rho$. Then,*

$$E[f^{(n)}(x)] > E[g^{(n)}(x)]$$

The proposition then follows. If the bidders acquire information symmetrically, obviously the expected maximum value of the n other bids increases since the standard deviation of $f_{\nu,\sqrt{1-\alpha^{2m}}\rho}$ increases as m increases from 1. As $m \to \infty$, the standard deviation approaches ρ and we have the situation when all bidders are informed. Thus, the expected maximum value is the greatest when all bidders are informed. □

In his paper, Parkes [11] solves the second-price sealed-bid auction equilibrium with uninformed agents using a symmetric equilibrium of the type described above. Every agent acquires the same amount of information. However, this leads to wasteful and unnecessary information acquisition amongst bidders with lower valuations, and insufficient information acquisition by bidders with higher valuations.

In an asymmetric Nash equilibrium, we want the right bidders to acquire the right amount of information. We want each bidder to acquire information optimally according to the maximum value implied by other bidders' optimal information acquisition strategies. Each bidder chooses to acquire information based on the expected "posted price" submitted by the highest bidder, i.e. $E[\theta_1]$. However, this maximum value is in turn determined by how much information each bidder chooses to acquire.

Nevertheless, in expectation we expect bidders with high values (close to $E[\theta_1]$) to acquire the most information and bidders with low values (much lower than $E[\theta_1]$) to acquire the least (if any) information. Unfortunately, this occurrence of asymmetric information acquisition does not fit into our Gaussian model of bidders' values, since the distribution becomes skewed. In particular, the rightmost segments of the distribution stretch out (as if it were the rightmost tail of $f_{\nu,\sqrt{1-\alpha^{2m}}\rho}$ for some $m > 1$), whereas the leftmost segments of the distribution don't move (as if no information is acquired and the values remain distributed according to $f_{\nu,\sqrt{1-\alpha^2}\rho}$). Even with asymmetric information acquisition, so long as the highest value bidders are still acquiring information, the expected maximum value of the n bids increases.

Although it is theoretically possible to find a fixed point m_f number of information acquisition steps such that $E[\theta_1]$ such that $\theta_1 \sim f^{(n)}_{\nu,\sqrt{1-\alpha^{2m_f}}}$ will lead all the bidders to acquire an *ex post* optimal amount of information once they see what the actual highest bid is, we leave such a derivation to future research.

Instead, we choose to approximate the optimal information acquisition strategy by finding the maximum number of information acquisitions $\hat{m}$ that any given bidder expects the highest value bidders to take. We then estimate the expected maximum value we are facing by using the distribution $f_{\nu,\sqrt{1-\alpha^{2\hat{m}}}\rho}$ to model the other bids. This is obviously an overestimate since the maximum amount of information acquisition occurs. The bidder then decides whether to acquire information using the same strategy outlined in the posted-price auction analysis, where the posted price is equal to the expected highest order statistic of $f_{\nu,\sqrt{1-\alpha^{2\hat{m}}}\rho}$. Because the posted price is overestimated, she will acquire too little information if her mean is less than the posted price, and too much information if her mean is above the posted price. Because of the construction of $E[\theta_1]$, we expect the former situation to occur much more frequently than the latter.

Recall that for any bidder with standard deviation σ about her true value, the *maximum* number of information acquisition steps she will take is defined by $m^*_\mu(\sigma)$ as defined in 3.2.4. Since all bidders enter the auction with the same standard deviation σ_0, we can calculate $m^*_\mu(\sigma_0)$. This becomes our $\hat{m}$.

Since bidders are given no dynamic information about the other bidders in a SPSB auction, the best each bidder can do when facing uninformed bidders is to approximate her highest competitors' *expected* number of information acquisition steps and deduce an *expected* value for the "posted price", i.e. highest maximum bid, she faces. Here we have given an approach for a bidder to approximate that expected value.

Remark 4.2.2. When*all bidders are uninformed* in a SPSB auction, the optimal information acquisition strategy becomes much harder to compute, as a bidder's *uncertainty about others' bidders actions is magnified*, even if she knows $g(p)$. Without that knowledge of $g(p)$, a bidder has no choice but to follow some suboptimal best-case or ad-hoc strategy.

4.3 Numerical Example Revisited

We again come back to our couple who is buying the home in a second-price sealed-bid auction. We can now determine their optimal information acquisition and bidding strategies. Recall that the couple has an initial belief distribution (μ_0, σ_0) about their value where $\mu_0 = 10$ and $\sigma_0 = 2.5$. Also, they can acquire each piece of information for a cost $c = 0.01$ and their computational effectiveness $1 - \alpha$ is equal to 0.3. The couple is told that there are $n = 10$ other bidders in the auction.

Other Bidders All Informed

In this scenario, the couple assumes that all the other 10 bidders are informed of their exact values.

Now assume that the couple has a belief about the distribution ζ from which the other bidders' values are drawn. Assume that this belief comes from knowledge of past auctions for homes of similar quality. In particular, they believe that ζ is Gaussian with mean $\nu = 8$ and standard deviation $\rho = \frac{\sigma_0}{\alpha} = 3.57$. From the equations above, we can write down the formula for the distribution g over the maximum bid p, according to the couple's beliefs, as

$$g(p) = \frac{10}{3.57\sqrt{2\pi}} e^{-\frac{(p-8)^2}{25.49}} \Phi\left(\frac{p-8}{3.57}\right)^9,$$

which implies

$$E[p] = 13.493.$$

The couple now proceeds as they would in a posted-price auction where $p = 13.493$. In particular, they find that $m^{**} = 1.13946$. Notice that $H(1.13946, 13.493) = 0.0106 > 0$, so there is no need to compute m^* since $m^* > m^{**}$ and $H(m^*, p) > H(m^{**}, p) > 0$. So the couple acquires information. Using a simulator, we find that $\mu_1 = 10.605$ and $\sigma_1 = \alpha\sigma_0 = 1.75$. Should the couple acquire information again? Plugging into the formula, we obtain $m^{**} = 1.379$. Since $H(1.379, 13.493) = -0.0045 < 0$, we calculate $m^* = 2.286$ and find that $H(2.286, 13.493) = -0.0026 < 0$. Thus the couple stops acquiring information.

Since $\mu_1 = 10.605$ and they are no longer acquiring information, the couple's optimal strategy is to bid truthfully $b^* = 9.8604$. Running a simulation and drawing $n = 10$ values independently from the distribution ζ, we obtain a maximum bid of 12.335. Thus, the couple's bid is not accepted and they do not buy the home. Overall, they lose \$0.01 from their one step of information acquisition.

Other Bidders All Uninformed

In this scenario, the couple assumes that all the other 10 bidders are uninformed of their exact values. Using the same distribution ζ over true values to be consistent, the couple believes that each of the other bidders i enters the auction with mean μ_0^i distributed according to a Gaussian f with mean $\nu = 8$ and standard deviation $\sqrt{(1-\alpha^2)}\rho = 2.55$, and standard deviation $\sigma_0 = \alpha\rho = 2.5$.

Using the approximation we describe above, we obtain the value for m^*, the maximum number of deliberations any given bidder will take and find that

$$m^* = 5.03.$$

So our couple approximates the final distribution they face as Gaussian with mean $\nu = 8$ and standard deviation $\sqrt{1-\alpha^{2m^*}}\rho = 3.52$. The expected maximum value of this distribution is calculated to be

$$E[p] = 13.416.$$

Notice that this is only slightly less than the $E[p]$ when all bidders are informed.

The couple now proceeds as they would in a posted-price auction where $p = 13.416$. In particular, they find that $m^{**} = 1.1097$. Since $H(1.1097, 13.416) = 0.0123 > 0$, the couple acquires information. Using the same simulation as above, we obtain that $\mu_1 = 10.605$ and $\sigma_1 = 1.75$. Should the couple continue acquiring information? They calculate $m^{**} = 1.338$ and $m^* = 2.286$, where $H(2.286, 13.416) = 0.00005 > 0$, and so they acquire information, finding that $\mu_2 = 9.8604$ and $\sigma_0 = \alpha\sigma_1 = 1.225$. Should they acquire yet again? Now, $m^{**} = 2.42089$. This time, we notice that $\chi(m^{**}) = 0.0002 < c = 0.01$ and so $H(m, 12.616)$ is decreasing in m. So there is no m such that $H(m, p)$ has positive value and so the couple stops acquiring information.

Since $\mu_2 = 9.8604$ and they are no longer acquiring information, the couple's optimal strategy is to bid truthfully $b^* = 9.8604$. To simulate the other uninformed bidders, we draw $n = 10$ values independently from a Gaussian distribution with $\nu = 8$ and standard deviation $\sqrt{(1-\alpha^2)}\rho = 2.14$. For each value, we calculate the optimal m^* number of acquisitions and have each bidder acquire information $\lfloor m^* \rfloor$ times. We then obtain a maximum bid of 11.539. Thus, the couple's bid is not accepted and they do not buy the home. Their overall loss is $2(0.01) = \$0.02$ from acquiring information two times.

Thus, in the case where we have all uninformed bidders, both the revenues to the seller as well as the gains to the bidder are decreased compared to the case where everybody is informed. We look into this effect in more detail in Chapter 6.

Chapter 5

Ascending-Price Auction

In this chapter, we formulate and solve the equilibrium bidding and information acquisition strategies for a bidder in an asynchronous ascending-price auction. We analyze the optimal bidding and information acquisition strategies in two cases: (a) when a bidder believes that all the other bidders are *informed* of their exact values and do not acquire information, and (b) when a bidder believes that all the other bidders are *uninformed* and can acquire information.

Like in the SPSB, in the ascending-price auction, whether a bidder wins the good depends on the bids of other bidders. However, bidders are given dynamic feedback about the other bids in the ascending-price because of the price level p. Namely, they not only have a belief distribution about the highest bid but they also know that the highest bid will be greater than p. Notice that this p is not like a posted price for the bidder because if she bids at p, she is not guaranteed to win the good at that price. Instead, p increases over time.

In this analysis, we make a simplifying assumption that a bidder can acquire at most a fixed $\tilde{m}$ pieces of information each time she decides to acquire information. As opposed to the posted-price and SPSB equilibrium information acquisition strategies where the bidder acquires information so long as $H(m, p) > 0$ or $J_p(m) > 0$ for some m, in this auction, we limit bidders that want to acquire information to acquiring between 1 and $\tilde{m}$ pieces of information. Since the auction must wait while bidders acquire information, this restriction on bidders is not unrealistic.

When hard valuation problems are introduced to the asynchronous ascending-price auction, the optimal bidding strategy becomes closely intertwined with the optimal information acquisition strategy. In fact, when many bidders can acquire information, an intricate model of interaction between bidders emerges. In particular, some bidders may become locked into a *waiting game* in which each bidder prefers to sit tight and wait for some other bidder to acquire information and to place a new bid that will increase the price. Perhaps the price will increase above any value the bidder would consider paying and she will never have to acquire information, saving her the costs.

5.1 One Uninformed Bidder

First, we consider the case with one uninformed bidder that can acquire information and that assumes all the other bidders are informed of their values.

5.1.1 Optimal Bidding Strategy

In the asynchronous ascending-price auction, at any given price level p, the bidder has three options: (1) bid, (2) remain active, or (3) drop out.

When a bidder is completely informed about her value, v, her optimal bidding strategy will be to submit bids until $p = v$. At this point, she will no longer bid. Once $p = v + \epsilon$ as a result of another bid, she will drop out. Because the bidder knows her exact value, she never chooses to remain active. Instead, informed bidders are forthcoming and want to push up the price to their threshold. They have no uncertainty and hence never induce a waiting game.

Now consider an uninformed bidder with beliefs (μ, σ) about her value. So long as the auction is still running, she does not *have to* bid or to drop out. Instead, she can just remain active, which always has nonnegative utility (because remaining active implicitly includes the option to drop out at any time). To keep the auction running smoothly, we assume that if a bidder is indifferent between bidding or remaining active, she will bid and that if a bidder is indifferent between dropping out or remaining active, she will drop out.

Intuitively, because of her uncertainty, a bidder will choose not to acquire information only when the price level is so far below her mean that it will never be worth it for her to acquire information at that price or when the price level is so far above her mean that she will never acquire information past that price. Otherwise, she will remain active. However, if the auction reaches the waiting game and the auctioneer approaches the bidder, she must either bid or drop out. In this case, for a bidder who is no longer acquiring information, if $\mu > p$, the expected utility of bidding is $\mu - p$ if she wins the good at p and 0 in all other cases, while the expected utility of dropping out is 0. Similarly, if $\mu < p$, the expected utility of bidding is negative, $\mu - p$, if she wins the good at p and 0 in all other cases, while the expected utility of dropping out is 0. Thus, after she is finished acquiring information, the bidder then bids if the price level is below her mean and drops out if the price level is above her mean.

It is helpful to define and characterize the following two *posted prices*.

Definition 5.1.1. Assume that the bidder is offered the opportunity to *buy* the good at price p and that the bidder must either bid or drop out. Then we define $p = \underline{p}$ as the price at which an uninformed bidder with beliefs (μ, σ) does not drop out (hence $\underline{p} < \mu$) and is indifferent between bidding and acquiring information.

Definition 5.1.2. Assume that the bidder is offered the opportunity to *buy* the good at price p and that the bidder must either bid or drop out. Then we define $p = \overline{p}$ as the price at which an uninformed bidder with beliefs (μ, σ) does not bid (hence $\overline{p} > \mu$) and is indifferent between acquiring information and dropping out.

We first consider $\underline{p}$. When facing a posted price $p < \mu$, a bidder either acquires information or buys the good at p. From Chapter 3, we know that the bidder's expected value for acquiring information at any posted price p is $H(m, p)$ where the bidder acquires m pieces of information. On the other hand, the bidder's expected value of bidding is $\mu - \underline{p}$. Thus, the condition for $\underline{p}$ is

$$H(\tilde{m}^*, \underline{p}) = \mu - \underline{p}, \tag{5.1.1}$$

where $\tilde{m}^* = \min(m^*, \tilde{m})$ and m^* is as derived in Chapter 3 for a posted price auction. If $H(\tilde{m}^*, \underline{p}) < 0$, then we say that $\underline{p} > 0$ does not exist. In our analysis, we assume $\underline{p} > 0$ exists.

Recall that $H(\tilde{m}, p)$ and $H(m^*, p)$ are increasing in p when $p < \mu$ (see Figure!!). So, when $\mu > p > \underline{p}$ the LHS of the equation increases whereas the RHS decreases. Thus, for p such that $\mu > p > \underline{p}$, the bidder prefers acquiring information. On the other hand, for p such that $\mu > \underline{p} > p$, the LHS of the equation decreases whereas the RHS increases, and thus we find that the bidder prefers to bid. So,

Lemma 14. *If p is the posted price at which a bidder with beliefs (μ, σ) is given the option to buy a good and $\underline{p}$ is defined as above, then $\underline{p} < \mu$ and (a) if $\mu > p > \underline{p}$, the bidder acquires information and (b) if $\mu > \underline{p} > p$, the bidder does not acquire information and bids.*

Now we consider $\overline{p}$. When facing a posted price $\underline{p} > \mu$, a bidder either acquires information or drops out. From Chapter 3, we know that the bidder's expected value for acquiring information at any posted price p is $H(m, p)$ where the bidder acquires m pieces of information. On the other hand, the bidder's value from dropping out is 0. Thus, the condition for $\overline{p}$ is

$$H(\tilde{m}^*, \overline{p}) = 0, \tag{5.1.2}$$

where again $\tilde{m}^* = \min(m^*, \tilde{m})$. If $H(\tilde{m}^*, \overline{p}) < 0$, then we say that $\overline{p} > 0$ does not exist. In our analysis, we assume that $\overline{p} > 0$ exists.

Recall that when $\mu < p$, both $H(\tilde{m}, p)$ and $H(m^*, p)$ are decreasing in p. Thus, for $\overline{p} > p > \mu$, the bidder prefers to acquire information since the LHS increases. On the other hand for $p > \overline{p} > \mu$, the LHS decreases and the bidder prefers to drop out. So,

Lemma 15. *If p is the posted price at which a bidder with beliefs (μ, σ) is given the option to buy a good and $\overline{p}$ is defined as above, then $\overline{p} > \mu$ and (a) if $\overline{p} > p > \mu$, the bidder acquires information and (b) if $p > \overline{p} > \mu$, the bidder does not acquire information and drop out.*

oNow we apply the same methodology as in the SPSB auction of considering a bidder's decision problem in light of a "posted price", namely her expected value for the maximum of the other bidders' values, i.e. $E[\theta_1]$. Taking the above two lemmas together, we have the following conclusion:

Proposition 16. *Let the posted price a bidder believes she is facing be denoted by $E[\theta_1]$. If $E[\theta_1] > \overline{p}$ or if $E[\theta_1] < \underline{p}$, she does not acquire information. Instead, she bids while $p < \mu$ and drops out when $p > \mu$. Otherwise, if $\overline{p} > E[\theta_1] > \underline{p}$, she will remain active.*

Two questions immediately arise: first, what is $E[\theta_1]$ in the asynchronous ascending-price auction? Secondly, what happens when the auction enters a waiting game and the bidder is approached by the auctioneer? She no longer has the choice of remaining active and must decide whether to acquire information and then whether to bid or to drop out.

Looking at $E[\theta_1]$

To answer the first question, we notice that, like in the SPSB auction analysis, we are able to use the posted price as a basic tool with which to analyze the ascending price auction. There are two main differences between SPSB and ascending-price auctions. First, a bidder is given feedback during the auction about other bids and thus dynamically updates the expected "posted price" she is facing. Each time a bid is placed, she is given new information about the other bids. Secondly, when new bids are placed, she is allowed to acquire more information in response to the new information, and so the ascending-price auctions allow for multiple stages of information acquisition.

At the beginning of the ascending-price auction, before bidding has begun, the bidder has the same beliefs about the other bidders' values as in the SPSB. We recall the function g from Chapter 4, which represents the bidder's belief distribution over the maximum value of the other bidders' values. So, at the beginning of the auction, the bidder believes the "posted price" she is facing is $E[\theta_1|\theta_1 \sim g]$. As the ascending-price auction proceeds, the bidder is given additional information about the maximum value, namely that $\theta_1 \geq p$ where p is the current price level.

In symbols, when the ascending-price auction is at price level p, the bidder believes that θ_1, the maximum value of the other bidders, is distributed according to $g(\theta_1|\theta_1 \geq p)$, where g is defined as in Chapter 4 and

$$g(\theta_1|\theta_1 \geq p) = \begin{cases} 0 & \theta_1 < p \\ \frac{g(\theta_1)}{\int_p^\infty g(\theta_1)\, d\theta_1} & \theta_1 \geq p \end{cases} \tag{5.1.3}$$

So $E[\theta_1] = E[\theta_1|\theta_1 \geq p]$, where $\theta_1 \sim g(\theta_1|\theta_1 \geq p)$. Notice that $E[\theta_1]$ is weakly increasing in p.

So far we have assumed secret dropout, but in fact, if we assume overt dropout, the bidder can also revise her expected value based on k, the number of agents left. This serves to give her more information with which to refine her beliefs about the maximum value she is facing. Now she can compute $g(\theta_1|\theta_k \geq p > \theta_{k+1})$, assuming that all the θ_i are drawn from a common distribution ζ, and then look at $E[\theta_1] = E[\theta_1|\theta_k \geq p > \theta_{k+1}]$. So overt dropout gives the bidder even more information about $E[\theta_1]$. For simplicity, however, we assume that we have secret dropout. The bidder does not know how many agents are left until the auctioneer comes up to her, at which point she knows there are only two bidders left, herself and the highest bidder.

Compared to the SPSB, as the auction progresses the bidder becomes much more well-informed of the highest bid she is facing and can choose not to acquire information until the price has reached a level where she finds it necessary. Thus, bidders are able to make better decisions about when to incur the cost and to acquire information, when just to bid, and when just to drop out.

A Brief Aside - p^* and p^{**}

We begin by defining the following two symbols to help keep our reasoning clear.

Definition 5.1.3. Let $\mathbf{p}^*$ be the price level in an ascending-price auction such that $E[\theta_1|\theta_1 > p^*] = \overline{p}$. So if $p > p^*$, a bidder will not acquire information and drops out when $p > \mu$.

Definition 5.1.4. Let $\mathbf{p}^{**}$ be the price level in an ascending-price auction such that $E[\theta_1|\theta_1 > p^{**}] = \underline{p}$. So if $p < p^{**}$, a bidder will never acquire information, bids while $p < \mu$, and drops out when $p > \mu$.

Obviously, $p^{**} < p^*$ since $E[\theta_1|\theta_1 > p]$ is increasing in p and $\overline{p} > \underline{p}$. Also, since p increases over time, once $p > p^*$, the bidder never acquires information and waits until $p > \mu$ to drop out.

We are interested in the relative positions of μ and p^* and p^{**}. Notice that intuitively p^* is the price level that implies an $E[\theta_1]$ so high that the bidder thinks she is not competitive and would rather not waste resources acquiring information. On the other hand p^{**} is the price level that implies an $E[\theta_1]$ so low that the bidder thinks she is easily the highest bidder and does not think she needs to acquire information. Thus it must be that $p^{**} < \mu_0$. We consider three cases.

1. $\mu_0 << E[\theta_1|\theta_1 > 0]$. In this case, the bidder may not acquire information even at the onset of the auction because the expected posted price she faces is too high. As the price level p increases, $E[\theta_1]$ only increases and so the bidder will never acquire information. In this case, p^{**} probably does not exist (there is never a price level where the bidder thinks she is easily the highest bidder). On the other hand, $\mu_0 >> p^*$ because right at the very beginning stages of the auction the bidder believes she is not competitive and will not want to waste her resources acquiring information. So we have $\mu_0 >> p^*$ *and* p^{**} *does not exist.* So even when $p > p^*$ the bidder will remain in the auction until $p = \mu_0$, but this does not have a big effect since by assumption μ_0 is very low.

2. $\mu_0 < E[\theta_1|\theta_1 > 0]$. In this case, p^{**} is either very low or nonexistent (there is only a very low price level (or no price level) where the bidder thinks she is easily the highest bidder). On the other hand, p may have to come relatively close to her mean μ before she believes that she is not competitive. Thus, we have $\mu_0 > p^* >> p^{**}$. Even when $p > p^*$, the bidder will remain in the auction until $p = \mu_0$, but again μ_0 is assumed not to be very high.

3. $\mu_0 \simeq E[\theta_1|\theta_1 > 0]$ or even $\mu_0 > E[\theta_1|\theta_1 > 0]$. In this case, at low price levels the bidder may believe that she is easily the highest bidder and does not need to acquire information. So p^{**}, although low, probably exists. On the other hand, because the bidder is told her value is likely to be the maximum bid, p will have to rise above her mean μ_0 before she will believe she is not competitive. So we have $p^* > \mu_0 > p^{**}$. So, once $p > p^*$, high bidders immediately drop out.

5.1.2 Optimal Information Acquisition Strategy

To understand the optimal strategies in the waiting game with one uninformed bidder, we begin by looking at how the auction progresses.

Let us assume the same setup as in Chapter 4 with $n + 1$ bidders indexed by $i \in \{0, 1, \ldots, n\}$. Let bidder 0 be the uninformed bidder with beliefs (μ, σ) about her true value. For all the other bidders, let their true values be θ_i for $i \in \{1, 2, \ldots, n\}$ such that $\theta_1 > \theta_2 > \ldots > \theta_n$.

Without loss of generality, we assume the price level p begins at 0 and recall that the bid increment is ϵ. Each informed bidder i follows her optimal bidding strategy of bidding until $p = \theta_i$ and then dropping out if a higher bid is placed. On the other hand, if $p < p^{**}$ or $p > p^*$, bidder 0 bids while $p < \mu$ and drops out if $p > \mu$; otherwise, if $p^* > p > p^{**}$ she remains active. The *waiting game* only occurs when bidder 1 is no longer willing to bid and bidder 0 is not bidding but remains active. To make our analysis more readable, we let the term "act according to her mean" refer to bidder 0 bidding if $p < \mu$ and dropping out if $p > \mu$. Also, let μ denote bidder 0's current mean after acquiring information, and let μ_0 denote bidder 0's initial mean before acquiring information.

In determining the optimal information acquisition strategy for bidder 0, we have the following cases to consider:

1. $\theta_1 > p^*$. In this case, bidder 1 continues bidding until either θ_1 or he is the only bidder left. When $p = \theta_2$, bidder 2 will stop bidding and drop out (and all other informed bidders other than bidder 1 will also have dropped out). Bidder 1 will continue driving up the price until θ_1. Since $\theta_1 > p^*$, during the process of bidder 1 driving up the price, bidder 0 will eventually act according to his mean and does not acquire information. If $\mu_0 > \theta_1$, bidder 0 wins the auction at price θ_1 without acquiring information. If $\mu_0 < \theta_1$, bidder 1 wins the auction at price $\max(\mu_0, \theta_1)$.

   Notice that in this case bidder 0 does not acquire information, and the waiting game does not occur.

2. $p^* > \theta_1$. We have two cases.

   (a) $p^{**} > \theta_1$. In this case, bidder 1 will continue bidding until $p = \theta_1$, after which he will not bid any further. At this point, the only bidders left will be bidders 0 and 1, since all the other informed bidders will have dropped out. Since $p^{**} > p$, bidder 0 will act according to his mean and not acquire information. If $\mu_0 > \theta_1$ bidder 0 wins the good at price θ_1. If $\mu_0 < \theta_1$ bidder 1 wins the good at price

μ_0.

(b) $\theta_1 > p^{**}$. In this case, bidder 1 will continue bidding until $p = \theta^{(1)}$, at which point she will not bid any further. At this point, the only bidders left will be bidders 0 and 1, since all the other informed bidders will have dropped out. Since $p^* > p > p^{**}$, bidder 0 will want to remain active but not to bid. So we are in the waiting game. At this point, the auctioneer will approach bidder 0 and demand that he either bid or drop out. This is the scenario we focus on when developing bidder 0's optimal information acquisition strategy. After acquiring information, if $\mu > \theta_1$, she will win the good at price θ_1; otherwise, bidder 1 wins at price $\theta_1 - \epsilon$.

Notice that the only situation in which bidder 0 acquires information is if $p^* > \theta_1 > p^{**}$.

If bidder 0 is the only uninformed bidder, when the auctioneer approaches her, she knows exactly which situation she is in. In particular, there is only one other bidder left, whose value is the current price level p. Thus, bidder n is facing exactly a posted-price! If she decides to bid at $p = \theta_1$, she will win the auction at θ_1.

To decide whether to bid or to drop out, the bidder follows exactly the information acquisition strategy outlined in Chapter 3, with a posted price equal to θ_1. The only caveat is that the bidder can acquire at most $\tilde{m}$ pieces of information. So, she acquires information if

$$H(\tilde{m}^*, \theta_1) > \max(0, \mu_0 - \theta_1).$$

The RHS of the equation denotes the bidder's action if she does not acquire information - she either bids or drops out, depending on the sign of $\mu_0 - \theta_1$.

Thus we have characterized the equilibrium information strategy when only one bidder is acquiring information.

Proposition 17. *In an asynchronous ascending-price auction with one uninformed bidder, if $\theta_1 > p^*$ or $\theta_1 < p^{**}$, the uninformed bidder acts according to her initial mean μ_0. If $p^* > \theta_1 > p^{**}$, the waiting game occurs at $p = \theta_1$. In this case, the bidder acquires information if and only if $H(\tilde{m}^*, \theta_1) > \max(0, \mu_0 - \theta_1)$.*

Remark 5.1.1. When there is only one uninformed bidder in the ascending-price auction, she does not need to know anything about $g(p)$! The auction mechanism gives her all the knowledge she needs about θ_1 (i.e. the value itself) and *it does not matter whether she* *is informed of either ζ, the distribution over the other bidders' values, or n, the number of other bidders.* In other words, the uninformed bidder's optimal information acquisition strategy is uninfluenced by her knowledge about the competition. This remark sounds familiar to Rezende's conclusion [14]; however, his model involved all uninformed bidders, not just one. The result we find here is in stark contrast to the situation with the SPSB auction.

5.2 All Uninformed Bidders

Now we turn our attention to the situation in which each bidder is uninformed and assumes that every other bidder is uninformed as well. So long as the auction is not in a waiting game, bidders will follow the optimal strategies outlined in the previous section. In particular, using the theory developed above, we have

Proposition 18. *While an ascending-price auction is running and has not reached a waiting game, an uninformed bidder acts according to her initial mean μ_0 if $p < p^{**}$ or if $p > p^*$, regardless of whether the other bidders are informed or uninformed. Otherwise, she remains active.*

Let there be n bidders in the auction. Let them be indexed with $i \in \{1, 2, \ldots, n\}$ such that $\mu_0^1 > \mu_0^2 > \ldots > \mu_0^n$, where μ_0^i is the mean of the initial belief distribution (μ_0^i, σ_0) with which bidder i enters the auction. According to our model, the μ_0^i are drawn independently from the same distribution γ (as defined in Chapter 2) and $\sigma_0 = \alpha\rho$ is the same for all bidders. As a result, we see that $p^{*1} > p^{*2} > \ldots > p^{*n}$ and $p^{**1} > p^{**2} > \ldots > p^{**n}$, where p^{*i} and p^{**i} denote p^* and p^{**} for bidder i respectively.

Assuming all the bidders follow the above strategy, the auction reaches a waiting game when $p = p^{**1}$. At this point, bidder 1 is assigned the good and all the bidders i such that $p^{*i} > p > p^{**i}$ are remaining active. Say there are k such active bidders. According to the auction protocol, the auctioneer then randomly selects one of the k active bidders, say bidder 2, and demands that she either bid or drop out. We assume that the auctioneer informs the bidders of k, the number of agents still left in the auction. What is bidder 2's optimal information acquisition strategy?

Unlike when there is only one uninformed bidder and so she knows exactly the posted price she faces, in this scenario the bidder has a more complicated situation. She knows there are k bidders left in the auction other than herself. Ideally, the bidder would rather remain active and have other people acquire information and bid, in hopes of the price rising past her p^* and making her drop out without incurring any costs. However, the auctioneer must force somebody to keep the auction running and has chosen bidder 2.

On first glance, this situation seems very similar to the situation at the beginning of the second-priced sealed-bid auction with k uninformed bidders. Bidder 2 must decide the estimated maximum value she is facing in order to determine whether to acquire information. She then follows her optimal bidding strategy. However, in this situation bidder 2 has several additional key pieces of information when she is approached, namely

1. the minimum above which the ultimate "posted price" for which the good will be sold lies, i.e. the current price level p,
2. the expected amount of information already acquired by each bidder still left in the auction and hence the current distribution over bidder's means, and
3. the maximum amount of information acquired by the highest bidders in the auction.

To see how she finds out this information, we look closely at the progress of the auction.

We assume that a bidder can see when the waiting game has begun and observes when other bidders are approached, when they bid, and when they drop out (i.e. as k decreases). Also, the auctioneer chooses a bidder *at random* to approach, so he chooses any given bidder with probability $\frac{1}{k}$. Thus, when a bidder is approached, she does not expect to be approached again for another k turns.

Let the waiting game begin at price level p'. Without loss of generality, say that bidder 2 is approached by the auctioneer at price level p'_2. She knows that (a) from observation s other bidders have been approached before her and that (b) $\frac{p'_2-p'}{\epsilon}$ of the s bidders have acquired information and chosen to bid instead of drop out. Notice that when a bidder acquires information in the waiting game she has no incentive to acquire information more than once because she would rather remain active and wait for somebody else to acquire information. Acquiring information once is the minimum she needs to appease the auctioneer and by doing so she just delays the rest of her information acquisition for later, when she will have even more information about the price level. So, at price level p' when k bidders are left, bidder 0 expects that that each bidder has acquired information $m' = \lfloor \frac{p'_1-p'}{k\epsilon} \rfloor$ times. Thus, each bidder is expected to have beliefs about her value with standard deviation $\alpha^{m'}\sigma_0$.

Given the standard deviation of a bidder's beliefs, recall our formula for the maximum value of $m^*_\mu(\sigma)$ in 3.2.4, for a bidder whose mean is equal to the price level and whose standard deviation is σ_0. Since the "posted price" that bidders expect to face is equal to the expected mean of the highest bidder, the maximum number of times that bidder 0 expects any bidder to acquire information in this waiting game before it ends is $m^*_\mu(\sigma)$! At any given time, she expects $\sigma = \alpha^m \sigma_0$ and so $m^*_\mu = m^*(\alpha^m \sigma_0)$.

In theory, if all the original bidders acquired information m^* times and none dropped out, their values would be distributed according to a Gaussian f with mean ν and standard deviation $\sqrt{1-\alpha^{2m^*}}\rho$. However, in reality we only have an altered version of that distribution, since only the highest value bidders have acquired m^* pieces of information, the middle value bidders have acquired less information, and the low value bidders have dropped out. A truly optimal equilibrium strategy would be for bidder 0 to use this altered distribution in her calculations.

To maintain the consistent properties of our model and analysis, however, we approximate that distribution with a Gaussian. In particular, we assume that the distribution of values over the n original bidders is exactly Gaussian f with mean ν and standard deviation $\sqrt{1-\alpha^{2m^*}}\rho$. This overestimates the value of the highest order statistic of the k values left because, as we found in Chapter 3, the expected maximum value increases in the standard deviation, holding the mean fixed. Still, we assume it is a reasonable approximation because what we are most interested in is what happens with the highest value bidders. More advanced statistics may be able to provide a method for modeling the distribution when a fraction of the bidders acquires information and the other fraction does not.

Thus, at any given price level p' when bidder 0 is approached by the auctioneer, she is equipped with an approximate distribution $f_{\nu,\sqrt{1-\alpha^{2m^*}}}$ over the values of the n bidders in the auction. She then calculates the expected value of the highest order statistic in that distribution, obtaining $E[\theta_1] = E[p]$. Now she faces a posted price! She then acquires information if there exists an $\tilde{m}' = \max(m^*, \tilde{m})$ such that $H(\tilde{m}', E[p]) > 0$. After acquiring information, she bids if $\mu_{\tilde{m}'} > p$ and drops out otherwise.

Remark 5.2.1. In an asynchronous ascending-price auction with uninformed bidders, the incentives are such that little wasteful information acquisition occurs. If a bidder knows her value is low relative to her beliefs about $E[\theta_1]$, she quickly drops out without acquiring information. If she knows her value is around the middle of the her expected distribution over θ_1, she remains in the auction for a little while but once p starts approaching her mean μ, she knows that the small chance that she is the high bidder is erased and she drops out. On the other hand, the high bidders know that they are high and stay in. As soon as one of their means drops below p, however, they drop out. When there are very few high-value bidders, this quickly becomes a problem as few bidders bid and many drop out. If this happens, the auction may be won at a price much lower than the the second highest bidder's value! This will hurt the seller's revenue.

What if the bidder does not know ζ or n? In this case, the only knowledge she has is that $\theta \geq p$ and that there are k other bidders in the auction. Thus, the best-case strategy she can adopt is to assume $\theta_1 = p$ and to respond accordingly. This is always an underestimate of the true maximum value over the bidders, and so if $\mu > \theta_1$ the bidder acquires too little information and if $\mu < \theta_1$, the bidder acquires too much. Nonetheless, the bidder still has some information about her competition in the auction and only has to acquire information (albeit occasionally wasteful) when the auctioneer approaches her. Again, this is in stark contrast to the situation in the SPSB auction where much wasteful deliberation occurs when uninformed bidders are not knowledgeable.

5.3 Numerical Example Revisited

We return to our couple. They now find themselves in an ascending-price auction for their dream home. The auctioneer tells them there are $n = 10$ other bidders.

All Informed Bidders

Like in the second-price sealed-bid auction, the couple assumes the 10 other bidders are informed and that their values are distributed according to ζ, a Gaussian with mean $\nu = 8$ and standard deviation $\rho = 3.57$. Because the auctioneer does not want to wait an excessively long time for any one bidder to acquire information, he imposes a condition that, if bidders decide to acquire information, bidders may acquire at most $\tilde{m} = 1$ pieces of information.

Recall the definitions above for $\underline{p}$ and $\overline{p}$. Solving these two formulas numerically,

$$\begin{aligned} H(1, \underline{p}) &= \mu - \underline{p} \text{ and} \\ H(1, \overline{p}) &= 0, \end{aligned}$$

we obtain

$$\begin{aligned} \underline{p} &= 9.515 \text{ and} \\ \overline{p} &= 13.842. \end{aligned}$$

So, according to their optimal strategy, the couple will bid at price level p if $E[\theta_1|\theta_1 > p] \leq 9.515$ and will drop out at p if $E[\theta_1|\theta_1 > p] \geq 13.842$, where θ_1 is the highest value of the other bidders.

What is $E[\theta_1|\theta_1 > p]$? From our analysis, we know

$$E[\theta_1|\theta_1 > p] = \int_0^\infty \theta_1 g(\theta_1|\theta_1 > p)\, d\theta_1 = \frac{\int_p^\infty \theta_1 g(\theta_1)\, d\theta_1}{\int_p^\infty g(\theta_1)\, d\theta_1}.$$

We want to find p^{**} such that $E[\theta_1|\theta_1 > p^{**}] = \underline{p} = 9.515$ and p^* such that $E[\theta_1|\theta_1 > p^*] = \overline{p} = 13.842$. Recall from our analysis of the SPSB numerical example, however, that $E[\theta_1|\theta_1 > 0] = 12.616$. Thus, for no $p > 0$ will $E[\theta_1|\theta_1 > p] = 9.515$. So there is no $p^{**} > 0$ and hence there is no price level for which the couple will bid rather than remain active until approached by the auctioneer. So the couple sits tight until the waiting game occurs. What about p^*? Solving numerically, we find that $E[\theta_1|\theta_1 > p^*] = 13.842$ gives us $p^* = 10.857$. So if the price level exceeds 10.857, the couple will not acquire information.

Running a simulation, we obtain values for the $n = 10$ other bidders and find that $\theta_3 = 10.352$, $\theta_2 = 11.026$ and $\theta_1 = 12.412$. So bidder 2 will bid until $p = 11.026$ and then drop out and bidder 1 will bid until $p = 12.412$ and then wait. At this point, if the couple is still in the auction, the auctioneer will approach the couple and require that they either bid or drop out. However, once $p \geq 10.857$, the couple will not acquire information because competition is too tough. Since $\mu = 10 < 10.857$, the couple drops out at 10.857. At this point there are two other bidders left, and bidding continues until $p = 11.026$, at which point bidder 1 wins the good at that price.

In this auction format, the couple leaves the auction with no net loss or gain because they do not acquire information. Compared to when their outcome in the SPSB auction, their utility is increased because in the ascending format, they have a lot of knowledge about θ_1, know that they are not competitive, and do not waste their resources in acquiring information.

All Uninformed Bidders

Now let's imagine that the couple now assumes that the other 10 bidders are in fact uninformed. Using the same distribution ζ over true values to be consistent, the couple believes that each of the other bidders i enters the auction with mean μ_0^i distributed according to γ with mean $\nu = 8$ and standard deviation $\sqrt{(1-\alpha^2)}\rho = 2.55$, and standard deviation $\sigma_0 = \alpha\rho = 2.5$. Let $\epsilon = 0.5$.

According to our analysis, the auction will run with bidders bidding until $p = p^{**1}$. Running a simulation, we find that $\mu_0^1 = 11.95$ and that there is no $p^{**1} > 0$. So the auction begins with no players bidding. On the other hand, even when nobody has bid, the maximum information acquisition steps anybody will take can be calculated to be

$$m^* = 5.03.$$

and so we obtain (keeping in mind that $\theta_1 > p$)

$$E[\theta_1] = 13.286.$$

With that expected price in the auction, six of the bidders drop out and four bidders remain in addition to the couple. So the auction begins in a waiting game. Let the other bidders be indexed $1, 2, 3, 4$.

We do not consider all the details of the auction, but instead consider only the couple's viewpoint. We know the waiting game starts at $p = 0$. Say that at $p = 1.5$, the couple is approached. Since $E[\theta_1|\theta_1 > 1.5] = 13.286$ and $H(1, E[\theta_1]) = 0.013 > 0$, the couple acquires information and now has $\mu_1 = 10.605$ as before and $\sigma_1 = 1.75$. Say the couple is approached again at $p = 2.5$. Since $E[\theta_1|\theta_1 > 2.5] = 13.2863$ and $H(1, E[\theta_1]) = -0.0029 < 0$, the couple doesn't acquire information but bids (since $10.605 > 2.5$). Say that the bidder is approached again at $p = 4, 5.5, 7, 8.5, 10 and 11$. At $p = 4, 5.5, 7, 8.5, 10$, $H(1, E[\theta_1]) < 0$ and so the couple does not acquire information but bids since $p < \mu = 10.605$. However, at $p = 11$, the couple calculates that $E[\theta_1] = 13.6724$ and does not acquire information. Since $p > \mu$, however, the couple decides to drop out. At this point there are two other bidders left.

The outcome of the auction for the couple is that their net loss is 0.01 from acquiring information once. This is an improvement over their outcome in the SPSB format with uninformed bidders. Nonetheless, the couple acquires more information than she needs to in the ascending-price format with informed bidders.

Chapter 6

Results and Discussion

In this chapter, we present the results of a basic comparison between the two dynamic auction mechanisms: the second-price sealed-bid and the asynchronous ascending-price. We do so by looking at the empirical implications of the optimal strategies we have derived.

We do not consider the posted-price because we believe that it depends too much on the part of the auctioneer to establish a revenue-maximizing price and any efficiency analysis will be unrealistic. When the bidders are informed, the auctioneer sets $p = \theta_1$ and always achieves 100% in revenues; when bidders are uninformed, the auctioneer sets $p = E[\theta_1]$ and either achieves 0 or $E[\theta_1]$ in revenues. Furthermore, any interesting information acquisition properties in the posted price auction, as we have shown, can be observed in an ascending-price auction with one uninformed bidder. The purpose of presenting an analysis of the optimal strategies in a posted-price auction has been to lay out a basic strategy for acquiring information given a posted price. As we have seen, using that strategy as a foundation is very useful in analyzing dynamic auctions.

Our main concern is the fact that the equivalence between the ascending-price auction and the second-price sealed-bid auction breaks down when bidders are not completely informed of their own private values, i.e. a hard valuation problem is introduced. In fact, we will see that each auction provides different knowledge to bidders and thus different incentives to incur costs and acquire information. Specifically, we find that the multi-stage nature of the ascending price auction provides much more knowledge to an uninformed bidder of what kind of competition she is facing.

When hard valuation problems are introduced and costs must be incurred for acquiring more information, we are primarily interested in whether ***the right bidders acquire the right amount of information***. In other words, we prefer auction mechanisms that provide incentives for the least competitive bidders (with low values) to acquire no or very little information, for the medium-value bidders to acquire just enough information to figure out they are not competitive, and for the highest-value bidders to acquire a significant amount of information to determine easily whether they in fact have the highest value. This point of concern is on behalf of the bidders' interests.

We are also concerned with *sellers' revenues.* Whether bidders decide to acquire info and bid is affected by their knowledge about competition. Since this knowledge varies between the two dynamic auctions, the resulting revenues may also differ. As Rezende points out [14], it is not entirely clear how information acquisition affects seller's revenues. If bidders begin the auction with abnormally high initial beliefs about their values (much higher than their true values), the seller will profit from those bidders not acquiring information. On the other hand, if bidders begin the auction with abnormally low initial beliefs, the seller will profit from their acquiring information.

Recall that there are two types of information that bidders may have: (a) knowledge about *the bidder's own private true value* and (b) information about *the highest value faced* by the bidder. The structure of the auction mechanism determines both *the incentives* for a bidder to acquire information about his own private value as well as how much knowledge (if any) bidders receive about other bidders' values. In particular, the more knowledge bidders receive about the highest bidder's value in the auction, the more "right" the amount of information they acquire.

6.1 One Uninformed Bidder

Consider the situation where only one bidder, bidder 0, is uninformed and the rest are informed. We compare the two auctions in (a) how much knowledge about θ_1 is provided and (b) what the incentives are to acquire information.

What knowledge about θ_1 is provided?

The second-price sealed-bid auction gives bidder 0 *no information* about θ_1. The only information she has is information we assume she has when she enters the auction. So she may or may not know ζ, the distribution from which bidders' true values are drawn, or n, the number of bidders in the auction.

On the other hand, the ascending-price auction with one uninformed bidder tells bidder 0 exactly the value of the highest bidder she faces. It does not matter whether the bidder knows ζ or n; she simply faces a posted price when the waiting game sets in. Thus, the structure of the asynchronous ascending-price auction fully informs a single uninformed bidder about her competition. The only uncertainty left is that in her own value.

How do incentives to acquire information compare?

Bidders' knowledge about θ_1 determines their incentives to acquire information. In the SPSB auction, the optimal information acquisition strategy is the same as that in the posted-price auction, using $E[p]$ as the posted price (assuming the bidder knows ζ and n). If $p = \theta_1$ is especially high or low, the SPSB auction does not alert the uninformed bidder in any way and thus she does not adjust her strategy.

In the ascending-price auction, the optimal information acquisition strategy is also the same as that in the posted-price auction, but this time using the actual value θ_1 as the posted price. This requires no knowledge of ζ or n. If θ_1 is especially high or low, the ascending-price auction alerts the uninformed bidder and she is able to adjust her strategy accordingly.

Empirical study

To test these conclusions empirically, we simulate several sets of $n = 20$ bidders, whose true values are distributed according to ζ with mean $\nu = 10$ and standard deviation $\rho = 3$. We initialize everybody at $\mu_0 = 10$ and then use a randomized function to help each bidder "acquire information." Repeating until the bidders' uncertainty decreases to 0, we obtain a set of values representing bidders' true values. This is by no means a comprehensive empirical study, but we have included it as a simple simulation in order to illustrate some major points and to encourage future empirical studies.

Our methodology is, for each set of bidders, to place ourselves as an uninformed bidder in the shoes of different bidders and to determine whether these bidders, if they were uninformed, acquire the "right" amount of information. We consider three situations: (a) the asynchronous ascending-price auction, (b) the second-price sealed-bid auction where the uninformed bidder knows ζ, n and hence $g(p)$, and (c) the second-price sealed-bid auction where the uninformed bidder knows neither ζ or n and hence has no information about the highest bid. We assume that in that last case the bidder uses a best-case strategy and assumes that the highest value is equal to her mean, $E[p] = \mu_0$. So she always acquires information $\lfloor m^*_\mu \rfloor$ times.

In our simulation we set $\alpha = 0.7$, $c = 0.01$, and $\sigma_0 = \rho\alpha = 2.1$. It then follows that

$\lfloor m^*_\mu \rfloor = 4$.

We consider the viewpoints of four different bidders: (a) a "low" bidder (the second lowest true value), (b) a "medium" bidder (the tenth lowest true value), (c) a "high" bidder (the third highest true value), and (d) the "highest" bidder (the highest value). We compare the performance of each auction according to the following measurements:

1. *number of steps of information acquisition*, m. The closer the highest maximum value is to a bidder's true value, the more she should acquire information.

2. *allocative efficiency*, η, where η is between 0 and 1 and is equal to the value of the bidder that wins the good divided by the value of the highest bidder in the auction.

3. *revenue*, R, where R is between 0 and 1 and is equal to the price paid by the bidder that wins the good divided by the value of the highest bidder in the auction (i.e. the maximum possible revenue).

4. *gain*, G, where G is the total gain of the uninformed bidder after the outcome of the auction.

When a bidder knows ζ, in the SPSB auction, she acquires information according to $E[p]$. In the ascending-price auction, she acquires information according to the exact $p = \theta_1$. So we consider three cases: (a) $p < E[p]$, (b) $p \simeq E[p]$, and (c) $p > E[p]$. The results of our simulations are in Tables 6.1 through 6.6.

Our results confirm our intuition about the key difference between the ascending-price and SPSB auctions. When the price is lower than expected, i.e. $p < E[p]$, the ascending-price auction alerts bidders of that fact and they adjust their information acquisition strategies accordingly. Across the board the low value bidder knows that her value is low and does not acquire information. If the uninformed bidder has a medium value, she adjusts and acquires some information to see if she might be competitive. If the uninformed bidder has a very high value, she realizes she is very competitive and acquires more information to determine whether she is in fact the highest bidder.On the other hand, in the second-price sealed-bid auction, the uninformed bidder has no idea that the maximum value is lower than expected, and so, if she has a medium value she assumes she is not competitive and does not acquire information. Even if the uninformed bidder actually has the highest value in the auction, it may be the case (as we see in our results) that her initial mean is low and she drops out. She is not aware that the values in general are low, and so loses the opportunity to win the good for a positive gain. Similarly, when the price is higher than expected, i.e. $p > E[p]$, the ascending-price auction also alerts bidders of that fact and, as we can see, the low and medium value bidders respond accordingly by not acquiring information. In fact, not even the high bidder acquires information in our simulation because she realizes that the maximum value is unexpectedly high. The SPSB auction, however, does not convey this knowledge and so the bidders do not react. Finally, when the price is around the expected value, the two auctions provide similar incentives for bidders to acquire information.

So, if the value of the highest bidder is much higher or lower than expected, the ascending-price auction informs bidders and provides the right incentives for the right bidders to acquire information about their private values. On the other hand, the second-price sealed-bid does not provide any information for bidders to change their information acquisition strategies. They always act as if the maximum value is equal to the expected value. Finally, notice that when the bidder does not have knowledge of $g(p)$, a lot of wasteful information acquisition occurs. Thus we have the following conclusion:

Result 19. *When there is only one uninformed bidder, the ascending-price auction lets the bidder know the highest value bid and thus provides incentives for the right bidders to acquire the right amount of information by adjusting based on the deviation of the actual highest bid from its expected value . The second-price sealed bid auction does not provide any knowledge about the maximum value and thus does not give incentives for bidders to adjust their information acquisition strategies. When the bidder does not know ζ, the second-price sealed-bid auction results in the uninformed bidder acquiring much more information than she should.*

We note that when there is only one uninformed bidder, the revenues of both auctions are quite similar. If the uninformed bidder has the highest value and does not acquire enough information, it is possible that an inefficient outcome results, but the revenues to the seller remain high because competition is tough. In each auction, the seller is guaranteed at least θ_2, the second highest value of the informed bidders.

6.2 All Uninformed Bidders

When all the bidders are uninformed, we have seen that it is important in the optimal information acquisition strategies for both auctions to know ζ, the distribution from which bidders' values are drawn.

As in the case where only one bidder is uninformed, the second-price sealed-bid auction does not give any extra information to a bidder other than that with which she enters the auction. In other words, she may know ζ and n, but she does not get any feedback on how much information each bidder acquires. The best she can do is estimate how much information each bidder acquires, use that information to estimate the distribution of bidders' final values after information acquisition and hence the expected value of the highest bidder, i.e. $E[\theta_1]$. In the approximate strategy detailed in Chapter 4, a bidder will overestimate $E[\theta_1]$. If $\mu_0 > E[\theta_1]$, the bidder acquires too much information and if $\mu_0 < E[\theta_1]$, the bidder acquires too little information. If, however, bidders do not know ζ and instead follow a best-case strategy as in the previous section, the non-competitive bidders will inevitably acquire too much information. Thus the SPSB with uninformed bidders does not provide any knowledge about θ_1.

Now consider the ascending-price auction. This auction mechanism gives bidders a relatively large amount of information compared to the SPSB. Most importantly, the bidder is told a minimum price (i.e. the current price level) above which θ_1 lies. Once the waiting game begins, the bidder also can see how much information has been acquired. Furthermore, if we assume overt dropout, the bidder knows how many bidders are left and can use the information $\theta_k > p > \theta_{k+1}$ in refining her beliefs about $E[\theta_1]$. As the auction progresses, the bidder can in fact obtain a very good estimate of the actual θ_1. Thus it is much easier in the ascending-price auction for bidders to know their relative competitiveness in the auction and to acquire the "right" amount of information.

What about the sellers' revenues? In the SPSB auction, if a bidder does not acquire (enough) information, the seller may profit if his initial belief is especially high and may lose if his initial belief is especially low.

In the ascending-price auction, recall our analysis of the relative positions of $p^*, p^{**}, and \mu_0$. If there are bidders whose initial values are very close to $E[\theta_1]$ or greater than $E[\theta_1]$, then they will stay in the auction and acquire information at a price level greater than their means. Thus, in this scenario it is possible for the seller to have increased revenues regardless of whether those initial beliefs are high or low. The reason is that high-value bidders *know* that they are very competitive and are willing to take a chance and acquire information, even when the price level exceeds their means.

Bidders whose initial values are lower than or much lower than $E[\theta_1]$ may drop out of the auction without acquiring information. But this situation affects the seller's revenues the same way as in the SPSB auction. Even though they do not acquire enough information, they may help or hinder the seller's revenues based on their initial beliefs about their true values.

Conclusion 20. *When all bidders are uninformed, the asynchronous ascending-price auction provides better incentives than the second-price sealed-bid auction for the right bidders* *to acquire the right amount of information. Furthermore, it is possible in the ascending-price auction to achieve higher revenues than in the SPSB auction, because bidders have enough* knowledge *about their relative competitiveness.*

A quick reference to the synchronous ascending-price auction

Here we make a brief intuitive comparison between the synchronous ascending-price auction, also known as a Japanese auction, and the two dynamic auctions we have studied in this paper. In the Japanese auction, bidders do not have the option of remaining active. By not dropping out of the auction, they are implicitly bidding. When all the bidders are informed and only one bidder (bidder 0) is uninformed, bidder 0 is not given as much information as she is in the asynchronous auction. The reason is that, at every price level p, she must decide whether to acquire information, to bid, or to drop out. She cannot just wait until the other bidders push the price up; she must bid at every price level. Thus, the synchronous ascending-price auction provides much less information to an uninformed bidder about θ_1. Nonetheless, it still provides *some*. In particular, an uninformed bidder knows that for any given price level p, $\theta_1 \geq p$, and with overt dropout, she also knows $\theta_k > p > \theta_{k+1}$.

The synchronous ascending-price auction provides some dynamic feedback to bidders about the maximum value θ_1 and thus has the advantage over the SPSB auction of reducing wasteful information acquisition when θ_1 is different from expected. Nonetheless, it does not provide *as much* feedback as the asynchronous ascending-price auction. Thus, the synchronous ascending-price auction may be a viable alternative mechanism to the SPSB auction in reducing wasteful information acquisition, but does not seem to hold much advantage over the asynchronous auction.

Bidder	AP_m	SBI_m	SBU_m	AP_G	SBI_G	SBU_G
Low	0	0	4	0	0	-0.04
Medium	1	0	4	-0.01	0	-0.04
High	2	2	4	-0.02	-0.02	-0.04
Highest	4	0	4	0.64	0.68	0.64

Table 6.1: $p = 13.750 < E[p] = 15.602$

Bidder	AP_η	SBI_η	SBU_η	AP_R	SBI_R	SBU_R
Low	1	1	1	0.95	0.95	0.95
Medium	1	1	1	0.95	0.95	0.95
High	1	1	1	0.95	0.95	0.95
Highest	1	0.95	1	0.95	0.91	0.95

Table 6.2: $p = 13.750 < E[p] = 15.602$

Bidder	AP_m	SBI_m	SBU_m	AP_G	SBI_G	SBU_G
Low	0	0	4	0	0	-0.04
Medium	0	0	4	0	0	-0.04
High	0	0	4	0	0	-0.04
Highest	3	7	4	0.86	0.82	0.85

Table 6.3: $p = 15.306 \simeq E[p] = 15.602$

Bidder	AP_η	SBI_η	SBU_η	AP_R	SBI_R	SBU_R
Low	1	1	1	0.94	0.94	0.94
Medium	1	1	1	0.94	0.94	0.94
High	1	1	1	0.94	0.94	0.94
Highest	1	1	1	0.94	0.94	0.94

Table 6.4: $p = 15.306 \simeq E[p] = 15.602$

Bidder	AP_m	SBI_m	SBU_m	AP_G	SBI_G	SBU_G
Low	0	0	4	0	0	-0.04
Medium	0	0	4	0	0	-0.04
High	0	5	4	0	-0.05	-0.04
Highest	2	2	4	2.57	2.57	2.55

Table 6.5: $p = 17.218 > E[p] = 15.602$

Bidder	AP_η	SBI_η	SBU_η	AP_R	SBI_R	SBU_R
Low	1	1	1	0.85	0.85	0.85
Medium	1	1	1	0.85	0.85	0.85
High	1	1	1	0.85	0.85	0.85
Highest	1	1	1	0.85	0.85	0.85

Table 6.6: $p = 17.218 > E[p] = 15.602$

Chapter 7

Conclusion

In this paper, we have taken a mathematical approach to formulating a consistent model for mid-auction multi-stage information acquisition for uninformed bidders in auctions. This improves on the stylized models used by Compte & Jehiel [4] and Rezende [14] and also on the simple uniform model introduced by Parkes [11]. Using this new model, we have formulated optimal bidding and information acquisition strategies for the three canonical auction formats: the posted-price, the second-price sealed-bid, and the ascending-price.

Our analysis confirms many of the results found in earlier work. In particular, we have confirmed Sandholm's [17] observation that a bidder can make better decisions about whether or not to perform further computation about the value of a good if it is well-informed about the other bidders' bids. The ascending-price auction format provides bidders with much more knowledge about the highest bid faced, especially when only one bidder is uninformed, and thus bidders acquire the more "right" amount of information than in the SPSB auction.

The major finding of Compte & Jehiel [4] is that some bidders that do not acquire information in the SPSB auction will acquire information in the ascending-price auction, *conditional on there being overt dropout.* When there is secret dropout, however, Compte & Jehiel claim that bidders that do not acquire information in the SPSB auction will also not acquire information in the ascending-price auction. Thus, they conclude that while the ascending-price auction is preferable when there is overt dropout, with secret dropout, the ascending-price auction is equivalent to the SPSB.

Rezende [14], on the other hand, finds a "remarkably simple formula" to determine the optimal moment for a bidder to acquire information in an ascending-price auction with secret dropout and argues that it is optimal regardless of how the other bidders behave. Using this analysis, he concludes that the dynamic auctions are preferable over one-shot auctions. Furthermore, Rezende claims that the assumption of overt dropout has no qualitative effect on his analysis.

From the results of this paper, we are able to reconcile those contradictory findings. Compte & Jehiel fail to take into account the improvement even with secret dropout that the ascending-price auction has over the SPSB because bidders that *do* acquire information in the SPSB but shouldn't have because their values are too low, no longer acquire information in the ascending-price auction. This is in line with our finding that the ascending-price auction induces the right amount of information acquisition on the part of the right bidders. Parkes [11] also comes to that conclusion.

In his paper, Rezende suggests that it may be true that sellers' revenues are higher in dynamic auctions. We too note a similar intuition. It will be interesting to test these hypotheses in a simulation with many uninformed bidders.

We conclude that the ascending-price auction is much more effective than the second-price sealed-bid auction at providing knowledge to bidders about other bidders' values. So the dynamic feedback provided in the ascending-price auction reduces the amount of wasteful computation performed by bidders. Bidders care about not wasting their information acquisition efforts and incurring those costs only when they have a good chance of winning the auction. Thus, bidders should prefer the ascending-price auction. In fact, we believe that the asynchronous ascending-price auction that we study should be chosen over the synchronous one.

When many bidders are uninformed, the ascending-price auction may even lead to higher revenues than the second-price sealed-bid. If this is true, then even the sellers may prefer to implement the dynamic feedback auction as opposed to the single-shot one. The sellers' main concern is the revenues the auctions produce.

Future Work

The results found in this paper hint at a variety of directions for future work relating to mid-auction information acquisition in auctions with bounded-rational bidders. We list a few here:

- A detailed comparison between the asynchronous ascending-price auction and the synchronous one in the case where many bidders are uninformed.
- Extension of our model to accommodate multi-object and perhaps combinatorial auctions.

- Conducting an empirical study of the approximate information acquisition strategies detailed in this paper when all the bidders are uninformed, and checking to see when and whether the revenue problem exists in the asynchronous ascending-price auction.
- Designing new auction mechanisms that are optimized for reducing wasteful information acquisition while maximizing efficiency and revenues.

We hope that this paper will encourage future research in those and other areas.

Appendix A

Miscellaneous Proofs

Proof of Proposition 1

The following is Chiburis' derivation [3] that proves (2.1.1). For all $x \in \mathbb{R}$,

$$\begin{aligned}
& \int_{-\infty}^{\infty} f_{\mu,\sqrt{1-\alpha^2}\sigma}(y) f_{y,\alpha\sigma}(x) dy \\
= & \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi(1-\alpha^2)}\sigma} e^{-(y-\mu)^2/2(1-\alpha^2)\sigma^2} \frac{1}{\sqrt{2\pi}\alpha\sigma} e^{-(x-y)^2/2\alpha^2\sigma^2} dy \\
= & \frac{1}{2\pi\sigma^2\alpha\sqrt{1-\alpha^2}} \int_{-\infty}^{\infty} e^{-(y-\mu)^2/2(1-\alpha^2)\sigma^2-(x-y)^2/2\alpha^2\sigma^2} dy \\
= & \frac{1}{2\pi\sigma^2\alpha\sqrt{1-\alpha^2}} \int_{-\infty}^{\infty} e^{-\frac{\alpha^2(y-\mu)^2+(1-\alpha^2)(x-y)^2}{2\alpha^2(1-\alpha^2)\sigma^2}} dy \\
= & \frac{1}{2\pi\sigma^2\alpha\sqrt{1-\alpha^2}} \int_{-\infty}^{\infty} e^{-\frac{\alpha^2(y^2-2y\mu+\mu^2)+(1-\alpha^2)(x^2-2xy+y^2)}{2\alpha^2(1-\alpha^2)\sigma^2}} dy
\end{aligned}$$

We now make the substitution $z = \frac{y-\alpha^2\mu-(1-\alpha^2)x}{\sigma\alpha\sqrt{1-\alpha^2}}$. Then $dz = \frac{dy}{\sigma\alpha\sqrt{1-\alpha^2}}$, and $z^2 =$

$\frac{y^2+\alpha^4\mu^2+(1-\alpha^2)^2x^2-2\alpha^2\mu y-2(1-\alpha^2)xy+2\alpha^2(1-\alpha^2)\mu x}{\sigma^2\alpha^2(1-\alpha^2)}$. Continuing, we get

$$
\begin{aligned}
&= \frac{1}{2\pi\sigma}\int_{-\infty}^{\infty} e^{-\frac{\sigma^2\alpha^2(1-\alpha)^2z^2+\alpha^2\mu^2+(1-\alpha^2)x^2-\alpha^4\mu^2-(1-\alpha^2)^2x^2-2\alpha^2(1-\alpha^2)\mu x}{2\alpha^2(1-\alpha^2)\sigma^2}}dz \\
&= \frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{\alpha^2\mu^2+(1-\alpha^2)x^2-\alpha^4\mu^2-(1-\alpha^2)^2x^2-2\alpha^2(1-\alpha^2)\mu x}{2\alpha^2(1-\alpha^2)\sigma^2}}\left(\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}e^{-\frac{z^2}{2}}dz\right) \\
&= \frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{\alpha^2\mu^2+(1-\alpha^2)x^2-\alpha^4\mu^2-(1-\alpha^2)^2x^2-2\alpha^2(1-\alpha^2)\mu x}{2\alpha^2(1-\alpha^2)\sigma^2}} \\
&= \frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{1}{2\sigma^2}\left(\frac{\mu^2}{1-\alpha^2}+\frac{x^2}{\alpha^2}-\frac{\alpha^2\mu^2}{1-\alpha^2}-\frac{(1-\alpha^2)x^2}{\alpha^2}-2\mu x\right)} \\
&= \frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{1}{2\sigma^2}(\mu^2+x^2-2\mu x)} \\
&= \frac{1}{\sqrt{2\pi}\sigma}e^{-\frac{(x-\mu)^2}{2\sigma^2}} \\
&= f_{\mu,\sigma}(x),
\end{aligned}
$$

since $\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}e^{-\frac{z^2}{2}}dz$ is an integral over a Gaussian probability density and therefore is equal to 1.

Proof of Proposition 2

The following is the inductive proof proposed by Chiburis [3] to prove (2.1.2). The inductive step relies on the fact that the convolution of two Gaussians with the same mean μ, but one with standard deviation x and one with standard deviation y is a Gaussian with mean μ and standard deviation $\sqrt{s^2+y^2}$ [6].

So, suppose we have deliberated $m-1$ times. By the inductive hypothesis, our current mean is chosen according to a Gaussian with mean μ and standard deviation $\sqrt{(}1-alpha^(2(m-1)))\sigma$, and our current beliefs have standard deviation $\sigma_{m-1}=\alpha^{m-1}\sigma$. In our mth deliberation, the new mean is chosen according to a Gaussian with standard deviation $\sigma_{m-1}\sqrt{1-\alpha^2}$. Then in the overall process, the new mean is chosen according to the convolution of a Gaussian with standard deviation $\sigma\sqrt{1-\alpha^{2(m-1)}}$ and a Gaussian with standard deviation $\sigma_{m-1}\sqrt{1-\alpha^2}$. By the convolution fact stated above, this is equivalent to choosing the new mean according to a Gaussian with standard deviation $\sqrt{\sigma^2(1-\alpha^{2(m-1)}+\alpha^{2(m-1)}(1-\alpha^2))}=\sigma\sqrt{1-\alpha^(2m)}$.

Proof of the movement of $\chi(m)$ when $\mu = p$

If $\mu = p$, then we have

$$\chi(m) = \frac{\sigma\alpha^{2m}\log(\frac{1}{\alpha})}{\sqrt{2\pi(1-\alpha^{2m})}}.$$

Consider as $m \to 0$. In this case, the denominator approaches 0 while the numerator approaches 1. Hence, the whole term explodes ($\to \infty$). Now consider as $m \to \infty$. In this case, the denominator approaches $\sqrt{2\pi}$ while the numerator approaches 0. Thus $\chi(m)$ approaches 0.

Solving for m^{**}

Recall (3.2.5), the expression for m^{**} that we need to solve:

$$\alpha^{4m} - \alpha^{2m}(3 + (\frac{p-\mu}{\sigma})^2) + 2 = 0.$$

Applying the quadratic formula, we find that

$$\alpha^{2m^{**}} = \frac{1}{2}\left(3 + (\frac{p-\mu}{\sigma})^2 \pm \sqrt{\left(3 + (\frac{p-\mu}{\sigma})^2\right)^2 - 8}\right).$$

Notice, however, that α^{2m} is less than 1. Thus it must be the case that the "-" holds. Taking logs of both sides, we get our result.

Proof of Incentive Compatibility of SPSB with Uninformed Bidders

To prove that it is optimal for an uninformed bidder with beliefs (μ, σ) about her true value to place a bid equal to μ in a SPSB auction, we first recall the formula (4.2.4) for a bidder's expected utility given a bid b:

$$\hat{U}(b) = \int_{p=0}^{b} g(p) \int_{-\infty}^{\infty} (v-p)f(v)\, dv\, dp.$$

To find the b that maximizes this equation, we differentiate with respect to b and set the expression equal to 0, obtaining

$$g(b^*) \int_{-\infty}^{\infty} (v-b^*)f(v)dv = 0.$$

The solution is to set $b^* = \mu$, *regardless of the distribution over outside bids, $g(p)$*. This shows that the auction is incentive compatible for uninformed bidders.

Bibliography

[1] B. Arnold, N. Balakrishnan, and H.N. Nagaraja, *A first course in order statistics*, Wiley, 1992.

[2] J. Berger, *Statistical decision theory and bayesian analysis*, 2 ed., Springer-Verlag, 1985.

[3] R. Chiburis, *Modeling hard valuation and deliberation in the assignment problem*, CS286r Final Project, Harvard University, 2002.

[4] O. Compte and P. Jehiel, *Auctions and information acquisition: Sealed-bid or dynamic formats?*, Technical Report, CERAS and UCL, 2000.

[5] R. Engelbrecht-Wiggans, *On a possible benefit to bid takers from using multi-stage auctions*, Management Science **34** (1998), no. 9, 1109–1120.

[6] W. Feller, *An introduction to probability theory and its applications*, 2 ed., vol. 2, Wiley, 1971.

[7] K. Larson and T. Sandholm, *Bargaining with limited computation: Deliberation equilibrium*, Artificial Intelligence **132** (2001a), no. 2, 183–217.

[8] ———, *Costly valuation computation in auctions*, Proc. Theoretical Aspects of Rationality and Knowledge VII (2001b), 169–182.

[9] T.K. Lee, *Competition and information acquisition in first price auctions*, Economic Letters **19** (1985), 129–132.

[10] D. Levin and J.L. Smith, *Equilibrium in auctions with entry*, American Economic Review **64** (1994), 585–599.

[11] D. Parkes, *Optimal auction design for agents with hard valuation problems - long version*, Proc. IJCAI'99 Workshop on Agent Mediated Electronic Commerce (1999).

[12] ———, *Auction design with costly preference elicitation*, Kluwer Academic Publishers, Netherlands, 2003.

[13] E. Rasmusen, *Strategic implications of uncertainty over one's own private value in auctions*, (2001), Indiana University.

[14] L. Rezende, *Mid-auction information acquisition*, Technical Report, Stanford University, Department of Economics, 2002.

[15] J. Rice, *Mathematical statistics and data analysis*, 2 ed., Duxbury Press, 1995.

[16] S. Russell and E. Wefald, *Principles of metareasoning*, Artificial Intelligence **49** (1991), 361–395.

[17] T.W. Sandholm, *Limitations of the vickrey auction in computational multiagent systems*, Second International Conference on Multiagent Systems (ICMAS-96) (1996), 299–306.

[18] W. Vickrey, *Counterspeculation, auctions and competitive sealed tenders*, Journal of Finance **16** (1961), no. 1, 8–37.