Chapter 1

Introduction

CONSIDER A seller of an object who faces $\mathcal{N}$ potential buyers. The seller may have some notion concerning the object's value to himself, but little or no information concerning how much any one of the potential buyers values the object. How should the seller choose to sell the object? A variety of different selling mechanisms exists. One commonly used mechanism involves announcing a take-it-or-leave-it price and then selling the object to the first person who accepts that price. Another might involve the seller's engaging in pair-wise negotiations with individual potential buyers, either simultaneously or sequentially. Yet a third way is to sell the object at auction.

Auctions are ubiquitous in market economies; they are also ancient, their durability suggesting that auctions serve an important allocational role. Over the past forty-five years, economic theorists have made considerable progress in understanding the factors influencing prices realized from goods sold at auction. For example, they have found that the seller's expected revenue depends on the auction format employed as well as the amount of competition, the information available to potential buyers, and the attitudes of bidders toward risk.

But what does holding an auction entail? The description of an *auction format* typically involves outlining the rules governing how the potential buyers must behave during the selling process; to wit, how bids must be tendered, who wins the auction, what the winner pays, and so forth. We shall introduce a number of different auction formats later in this book. Most importantly, however, the seller must commit to abide by the rules under a particular auction format.

Perhaps the most important feature defining environments in which auctions are used involves the existence of an asymmetry of information between the seller and the potential buyers. Typically, the seller knows little or nothing concerning the valuations of potential buyers. Moreover, these potential buyers have no incentive to tell the seller anything about their valuations. The role of the auction format is to get the potential buyers to reveal to the seller information concerning their valuations of the object.

But how do the valuations of potential buyers obtain? The way in which potential buyers form their valuations remains an open question in economics. In fact, in auction theory, researchers are unusually vague concerning what generates the demand structure, unlike in standard demand theory where considerable care is taken to specify the structure of preferences. Suffice it to say that, in auction theory, when economic theorists come to modeling this asymmetry in information as well as the heterogeneity in valuations across agents, they use random variables. Often, it is assumed that each potential bidder demands at most one unit of the object in question. In the simplest model, the marginal utility of this one unit, for each potential bidder, is assumed an independent and identically distributed realization of a continuous random variable V which has a differentiable cumulative distribution function $F_V(v)$ and probability density function $f_V(v)$ equal to $dF_V(v)/dv$. By and large, the budget constraint as well as issues of substitution are ignored.

1.1 An Example

For example, in the most common paradigm of an auction, referred to in the preface as the independent private-values paradigm (IPVP), each of $\mathcal{N}$ potential bidders gets an independent and identically distributed draw $\{v_i\}_{i=1}^{\mathcal{N}}$ from $F_V(v)$. If one orders these $\mathcal{N}$ valuations

$$v_{(1:\mathcal{N})} \geq v_{(2:\mathcal{N})} \geq \dots \geq v_{(\mathcal{N}:\mathcal{N})}$$

and then plots the highest valuation first, for which aggregate demand at that price is one, and then the second-highest valuation next, for which aggregate demand at that price is two, and so forth, one obtains the step function of aggregate demand, which is depicted in figure 1.1 for $\mathcal{N}$ equal five.

One way to interpret $F_V(v)$ is as follows: First, define the *survivor function*

$$\Pr(V > v) = S_V(v) = [1 - F_V(v)] = [1 - \Pr(V \leq v)]$$

which is the proportion of the population having demand when the price is v. Plotting the price p on the ordinate, as economists are wont to do, and $\mathcal{N}S_V(p)$ on the abscissa, one has an *expected-demand curve* as is depicted in figure 1.2. Each potential bidder, of which there are $\mathcal{N}$, is assumed to demand at most one unit, so aggregate demand is at most $\mathcal{N}$ when p is zero.

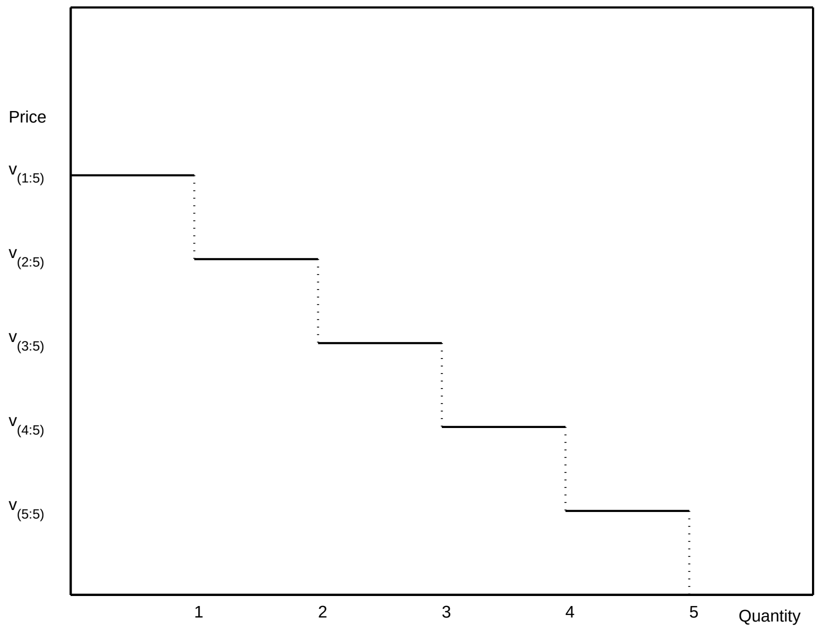

Figure 1.1
Aggregate Demand Step-Function

Now, the ordered marginal utilities of the $\mathcal{N}$ potential bidders are $\{v_{(i:\mathcal{N})}\}_{i=1}^{\mathcal{N}}$. From these one can create the sample analogue of $S_V(v)$, the empirical survivor function $\hat{S}_V(v)$; e.g., using the Kaplan–Meier product-limit estimator. We have depicted an estimate of expected demand, based on the estimated survivor function, generated from a sample of size $\mathcal{N}$ equal five, along with the population expected-demand function when $\mathcal{N}$ is five, in figure 1.3.

The idea of empirical work involving auction data is to estimate the expected-demand function $\mathcal{N}S_V(p)$ using the bids of the n participants at the auction. What makes this endeavor sometimes difficult, but invariably interesting, is that the n participants are often a subset of the potential bidders. Sample selection, in the Heckman sense, often exists. Also, depending on the auction format, bidders do not always reveal their true marginal utility.

In fact, one way to view auction theory is as demand analysis with a small number of consumers. But, unlike in standard demand analysis, where one typically assumes that the prices faced by an individual consumer are fixed, in auction theory one must take into account that the format and the rules of the auction, the primitive information giving

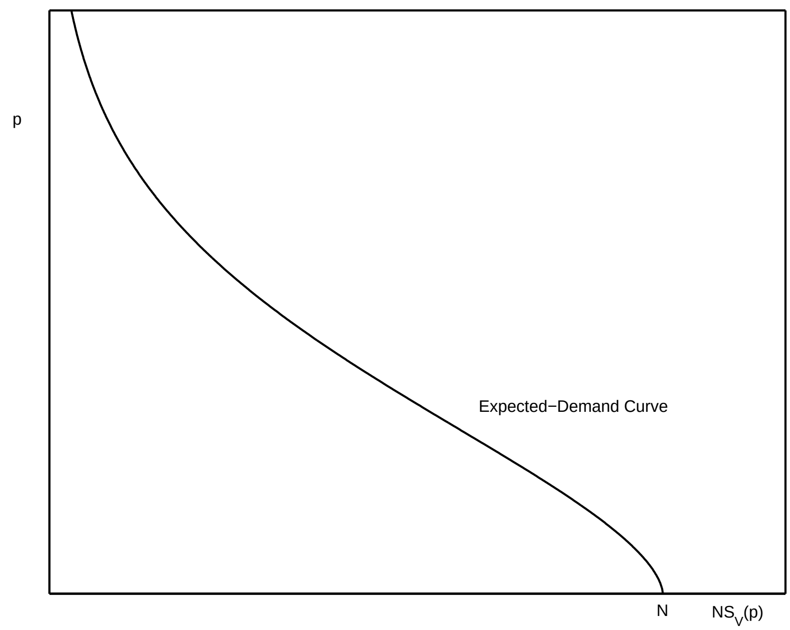

Figure 1.2
Expected-Demand Curve

rise to heterogeneity in beliefs concerning valuations, the preferences of potential bidders, the strategic behavior of the participants as well as the notion of equilibrium will *all* have an effect on the traded price, the winning bid. Thus, in the language of the econometrician, prices are *endogenous*. How can a researcher learn about the preferences of agents using either the bids submitted at auctions or just the winning bids?

The SEA uses the twin hypotheses of optimizing behavior and market equilibrium (henceforth optimization and equilibrium) to identify $F_V(v)$, the distribution of valuations.[1]

This is important. For, prior to applications of the SEA, many believed that it was impractical to implement mechanism-design theory to calculate the optimal selling mechanism because the optimal selling mechanism depended on quantities typically unobserved by the designer; viz., $F_V(v)$. Moreover, the actions (equilibrium strategies) of the agents, their bids $\{s_i\}_{i=1}^n$, while positively related to the valuations, were not always fully revealing; e.g., specifically, in the case of first-price, sealed-

[1] Typically, one assumes that potential bidders maximize the expected profit or the expected utility of profit from winning the auction and then uses either dominance or Bayes–Nash as an equilibrium concept.

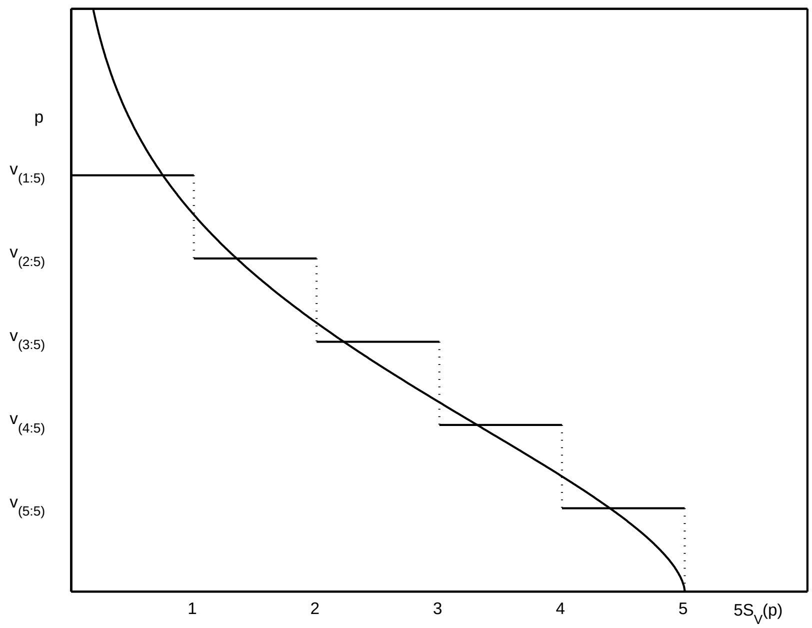

Figure 1.3
Empirical and Population Expected-Demand Curves

bid and oral, descending-price auctions. Thus, some believed that no way existed to estimate this distribution of latent types $F_V(v)$.

Now, one can typically estimate consistently the cumulative distribution function of observed actions (strategies) $F_S(s)$ using well-known empirical methods. Note, too, that in auction theory the strategy S is often a continuous and differentiable function σ of V. For example, at first-price, sealed-bid auctions the Bayes–Nash, equilibrium-bid function is

$$\sigma(V) = v - \frac{\int_0^v F_V(u)^{\mathcal{N}-1}\, du}{F_V(v)^{\mathcal{N}-1}}. \tag{1.1}$$

In general, when

$$S = \sigma(V) \quad \sigma'(V) > 0,$$

as is the case of (1.1),

$$V = \sigma^{-1}(S)$$

and

$$f_S(s) = \frac{f_V[\sigma^{-1}(s)]}{\sigma'[\sigma^{-1}(s)]}.$$

Thus, under suitable regularity conditions, which are usually met in theoretical models of auctions, one can construct an estimator $\hat{F}_S(s)$ of $F_S^0(s)$ where the superscript "0" denotes the true population value. Moreover, as the sample size increases, the estimator $\hat{F}_S(s)$ converges in probability to $F_S^0(s)$. Also, from $\hat{F}_S(s)$ one can usually construct a consistent estimator $\hat{f}_S(s)$ of $f_S^0(s)$ and subsequently, $\hat{F}_V(v)$, a consistent estimator of $F_V^0(v)$. Thus, the SEAD is an econometric *identification strategy*.

One can quite rightfully ask why, given the detailed research existing in standard demand theory, would one use such a blunt instrument to investigate demand? In the SEAD, strategic behavior is the most important consideration. Reverse engineering in the face of deception by market participants is the goal. Thus, all other considerations, typically deemed important in standard demand theory, have been shunted to the side in order that the main focus not be lost.

1.2 Some Intriguing Problems

One of the most intriguing problems faced by researchers who investigate data from auctions using the SEA is that different auction formats typically generate different kinds of information, so one omnibus empirical procedure to analyze these data cannot be proposed. One can, however, propose a general strategy; in this book, we describe several of many recent contributions to this general strategy.

To see how different auction formats generate different amounts of information, consider first the most informative auction format, the second-price, sealed-bid auction. At second-price, sealed-bid auctions within the IPVP, as will be shown below, each of the $\mathcal{N}$ bidders reveals his valuation truthfully. Thus, in the absence of a minimum bid price, the empirical distribution of bids $\{b_i\}_{i=1}^{\mathcal{N}}$ can be used to estimate the cumulative distribution of valuations. To wit, construct $\hat{F}_V(v)$, an estimate of $F_V^0(v)$, using the empirical distribution function

$$\hat{F}_V(v) = \frac{1}{\mathcal{N}} \sum_{i=1}^{\mathcal{N}} \mathbf{1}(b_i \le v)$$

where $\mathbf{1}(\mathsf{A})$ denotes the indicator function of the event A. The identifying assumption in this case is that bidders tell the truth, bid their actual valuations

$$B_i = \beta(V_i) = V_i.$$

Of course, the properties of $\hat{F}_V(v)$ can be improved by kernel-smoothing; in other words, using the estimator

$$\tilde{F}_V(v) = \frac{1}{\mathcal{N}} \sum_{i=1}^{\mathcal{N}} K\left(\frac{b_i - v}{h_F}\right)$$

where h_F is often referred to as the *bandwidth parameter*, while the function $K(\cdot)$, which is often referred to as the *cumulative kernel function*, has the following properties:

$$\begin{aligned} \lim_{y \to -\infty} K(y) &= 0 \\ \lim_{y \to \infty} K(y) &= 1 \\ K(y) &\geq 0 \quad -\infty < y < \infty. \end{aligned}$$

Unfortunately, the standard asymptotics are typically undertaken as $\mathcal{N}$ goes to ∞ and this does not happen at an auction. What to do?

Typically, to get more data, researchers combine data from auctions of objects that are not exactly alike. Thus, the independent and identically distributed assumption commonly made in empirical work may not apply. In some cases, the objects for sale may differ in observable ways that can be summarized for auction t by an observed vector of covariates $\boldsymbol{z}_t$. If one is willing to adopt a single-index model, then one can write

$$V_{it} = \mu(\boldsymbol{z}_t^{\top} \boldsymbol{\gamma}) + U_{it}$$

where

$$\mathcal{E}(U_{it} | \boldsymbol{z}_t) = 0,$$

and use the methods discussed in Horowitz (1998). Of more concern than the dearth of data concerning identical objects is the fact that second-price, sealed-bid auctions are rarely, if ever, used.

The most commonly used auction format and, under certain assumptions to be outlined below, also the next most informative format, from the perspective of an econometrician, is the oral, ascending-price auction. As we shall see below, at oral, ascending-price auctions within the IPVP, assuming the *clock model* of Milgrom and Weber (1982), each nonwinning bidder reveals his valuation truthfully, while all one knows about the winner is that his valuation is above the second-highest valuation. It is in this last sense that data from oral, ascending-price auctions are not as informative as second-price, sealed-bid auctions. Now, under clock model assumptions, the cumulative distribution function $F_W(w)$

of the winning price W is the cumulative distribution function of the second-highest valuation $V_{(2:\mathcal{N})}$ which is defined by

$$F_W(w) = \mathcal{N}(\mathcal{N}-1) \int_0^{F_V(w)} u^{\mathcal{N}-2}(1-u)\ du.$$

In the absence of a minimum bid price, when the number of potential bidders $\mathcal{N}$ is fixed and known, the empirical distribution of winning bids $\{w_t\}_{t=1}^T$ for a sample of T auctions, can be used to estimate the cumulative distribution of valuations by solving the following equation:

$$\begin{aligned}\hat{F}_W(v) &= \frac{1}{T}\sum_{t=1}^{T} \mathbf{1}(w_t \leq v) \\ &= \mathcal{N}(\mathcal{N}-1) \int_0^{\hat{F}_V(v)} u^{\mathcal{N}-2}(1-u)\ du\end{aligned}$$

at each point v. In this case, the identifying assumptions are that bidders tell the truth

$$B_i = \beta(V_i) = V_i$$

and that the winning price is the second order statistic of valuations

$$W_t = \beta[V_{(2:\mathcal{N}),t}] = V_{(2:\mathcal{N}),t}.$$

Of course, kernel-smoothing methods can improve the small-sample behavior of the estimator. Also, kernel-smoothing methods are needed to provide estimates of the optimal selling mechanism, which takes the form of an optimal minimum bid price ρ^* solving the following equation:

$$\rho^* = v_0 + \frac{[1 - F_V(\rho^*)]}{f_V(\rho^*)}$$

where v_0 is the seller's valuation of the object for sale.

Another intriguing problem arises when a binding minimum bid price exists because, in that circumstance, not all of the $\mathcal{N}$ potential bidders may participate at the auction. For example, when the minimum bid price is r, the number of *participating* bidders N is a random variable defined by

$$N = \sum_{i=1}^{\mathcal{N}} \mathbf{1}(V_i \geq r).$$

Because each of the random variables $\{\mathbf{1}(V_i \geq r)\}_{i=1}^{\mathcal{N}}$ is an independent and identically distributed Bernoulli random variable, their sum, the random variable N, is distributed binomially with two parameters, $\mathcal{N}$

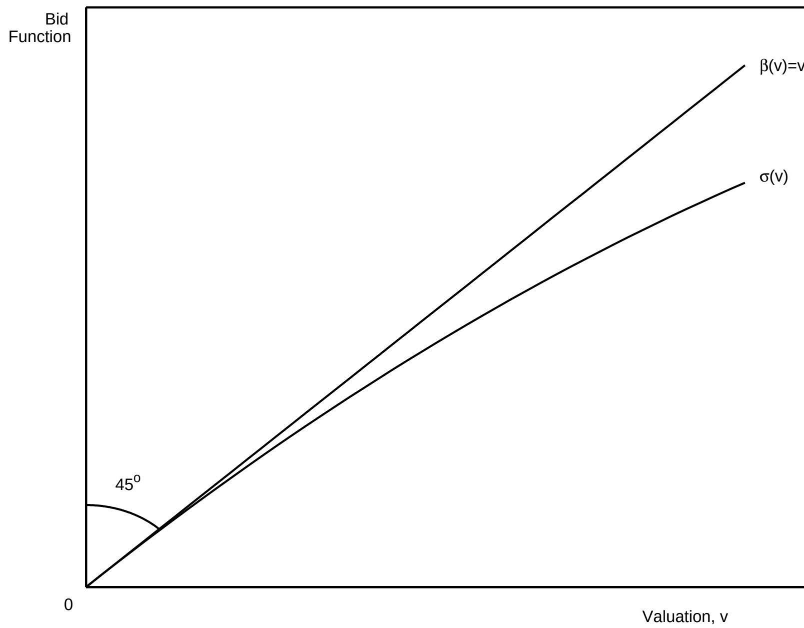

Figure 1.4
Graphs of Bid Functions:
First-Price and Second-Price, Sealed-Bid Auctions

and $[1 - F_V(r)]$. Moreover, the bids submitted at auction represent a *truncated* sample: only those potential bidders having valuations greater than r appear in the sample.

One feature that makes oral, ascending-price and second-price, sealed-bid auctions particularly tractable, at least numerically, is that the bid function in each case is a trivial function of the valuation. This is not the case at either first-price, sealed-bid or oral, descending-price auctions. At these auctions, within the symmetric IPVP and assuming risk-neutral potential bidders, the Bayes–Nash, equilibrium-bid function is

$$\sigma(v) = v - \frac{\int_0^v F_V(u)^{\mathcal{N}-1}\ du}{F_V(v)^{\mathcal{N}-1}}.$$

The first thing to notice about this strategy function is that it is monotonic, having positive slope less than one when v exceeds zero. What this means is that bidders with higher valuations bid more, but they bid systematically less than their true value. Moreover, the higher a potential bidder's valuation, the larger is the extent of this deception. Thus, the winner of either a first-price, sealed-bid or an oral, descending-price sealed-bid auction, the potential bidder with the highest valuation, will be the most deceptive: for a set of potential bidders, the winning bid will be the furthest from the actual valuation. We depict an example of this equilibrium-bid function at first-price, sealed-bid auctions in figure 1.4. In this figure, the 45°-line, denoted $\beta(v)$, depicts the equilibrium-bid function at the second-price, sealed-bid auction, while the curve that is everywhere below it (except at zero where they are equal), denoted $\sigma(v)$, depicts the equilibrium-bid function at the first-price, sealed-bid auction. Notice how a small change in the rules of the auction has an important impact on the behavior of the bidders at the auction.

The least informative format is the oral, descending-price auction. At these auctions, the researcher only gets to observe an action of the winner. A major portion of the potential bidders, $[(\mathcal{N}-1)/\mathcal{N}]$ to be exact, reveals no information. It is, in fact, somewhat surprising that one can make any statements concerning the cumulative distribution function $F_V(v)$ by just observing the winning bids from a sample of these auctions, particularly if a binding minimum bid price exists because the number of participants is then endogenous, but such is the power of the twin identifying hypotheses, optimization and equilibrium.

1.3 Plan of Book

In this book, we present an introduction to modern econometric techniques that are used in conjunction with the SEA to interpret field data from auctions. We do not consider the application of structural econometric methods to experimental data because, in those cases, the researcher typically knows $F_V^0(v)$ as he or she has selected it to generate the data for the subjects of the experiments. However, the methods we describe below can be used to test particular hypotheses using data from experiments. Thus, our book should be of interest to experimental workers who study auctions.

In chapter 2, we present an overview of single-object auction theory assuming auctions can be modeled as noncooperative games of incomplete information. We begin by describing the four most commonly studied auction formats and then some additional rules. Subsequently, we describe three models of information structures. Because the preferences of bidders are closely related to the structure of information, we discuss them next, but separately, to highlight the importance of risk aversion in formulating the decision problem faced by potential bidders under two of the auction formats. Ultimately, we derive the equilibrium-bid functions under the four auction formats, for risk-neutral and risk-averse bidders, without and with binding minimum bid prices. We then discuss the revenue equivalence proposition and characterize the optimal auction. We also show how risk aversion affects expected revenues under the four auction formats and outline Myerson's (1981) method for constructing optimal auctions. At the end of the chapter, we present a brief description of the winner's curse, perhaps the most well-known phenomenon in auction theory, showing why it is irrelevant for the models considered below.

While Krishna (2002) has provided an elegant and complete treatment of the material presented in chapter 2, we use this chapter to develop a notation, to introduce well-known results, and to outline the material necessary to formulate the major questions of interest.

In chapter 3, we then investigate the econometrics of oral, ascending-price and second-price, sealed-bid auctions. Even though it is rarely used, we begin with the second-price, sealed-bid auction because this format allows us to develop the basic intuition of the SEA within a nonparametric framework. Subsequently, we introduce covariates and then discuss single-index models and semiparametric estimation methods. The implications of a binding minimum bid price are discussed next. Endogenous participation, induced by a binding reserve price, highlights the limitations of nonparametric methods, so we introduce parametric methods.

Having developed most of the important econometric results concerning second-price auctions within the symmetric IPVP, we then present an extended policy application by Paarsch (1997), who estimated the optimal selling mechanism for timber in the province of British Columbia, Canada.

The fact that most oral, ascending-price auctions have either known bid increments or random, bidder-induced jumps in the price leads us naturally to an analysis of incomplete data and inference following the work of Haile and Tamer (2003).

In the final section of the chapter, we introduce asymmetric bidders, outlining a proof of nonparametric identification based on Meilijson (1981) and used by Brendstrup and Paarsch (forthcoming). The model as well as the methods of identification and estimation introduced in this section will prove useful in the specification of multi-unit auctions described in chapter 5.

In chapter 4, we investigate the econometrics of first-price, sealed-bid and oral, descending-price auctions. Following Paarsch (1989, 1992), we first derive the data-generating processes of the equilibrium-bid function as well as the winning bid. We then use the work of Guerre, Perrigne, and Vuong (2000) to demonstrate nonparametric identification. In the following section, we describe four different estimation strategies: First, we describe the nonparametric estimation methods of Guerre et al.

Subsequently, in an effort to deal effectively with observed covariate heterogeneity, we introduce parametric models, specifically discussing the method of maximum likelihood of Donald and Paarsch (1993, 1996) and then the method of simulated nonlinear least-squares of Laffont, Ossard, and Vuong (1995). Finally, we address some criticisms of the maximum-likelihood approach, examining the work of Donald and Paarsch (2002).

We then introduce a binding reserve price. In these cases, as noted by Brendstrup and Paarsch (2003), the extensive-form games at first-price, sealed-bid and oral, descending-price auctions are different because the number of participants is typically observed at oral, descending-price auctions, but realistically assumed unknown at first-price, sealed-bid auctions.

In most structural-econometric analyses of auction data, researchers have typically assumed that the potential bidders are risk neutral with respect to winning the auction. At oral, ascending-price and second-price, sealed-bid auctions, such an assumption is irrelevant because the dominant-strategy, equilibrium-bid function remains unchanged under these formats when potential bidders are risk averse. For example, at second-price, sealed-bid auctions, risk-averse bidders continue to bid their valuations when they exceed the minimum bid price. On the other hand, at first-price, sealed-bid and oral, descending-price auctions the attitudes of potential bidders toward risk matter. Thus, in the next section of chapter 4, we describe the effects of symmetric, von Neumann–Morgenstern preferences on the structural-econometric analysis, first using parametric methods, as in Donald and Paarsch (1996), and then using semiparametric methods, following Campo, Guerre, Perrigne, and Vuong (2000).

All of the surveyed research concerning first-price, sealed-bid and oral, ascending-price auctions has been within the symmetric IPVP. We go on to examine the effects of stochastic private-values, following the research of Lu (2004), who used the theoretical work of Ëso and White (2004). We also consider asymmetric bidders, those whose valuations are drawn from different distributions, especially in the presence of a binding reserve price, examining Brendstrup and Paarsch (2003), who have extended the results of Guerre et al. We then consider the work of Krasnokutskaya (2004), who investigated the effects of unobserved heterogeneity within the IPVP.

To illustrate a policy experiment, we examine the research of Brendstrup and Paarsch (forthcoming) in which the performance of the oral, ascending-price vis-a-vis the oral, descending-price auction is compared when potential bidders are asymmetric. Under these conditions, Maskin and Riley (2000) have demonstrated that inefficient allocations can obtain at oral, descending-price auctions, while at oral ascending-price auctions efficient allocations always obtain, at least within the IPVP. Moreover, the revenue equivalence proposition breaks down. Using data from fish auctions in Grenå, Denmark, Brendstrup and Paarsch estimated the incidence and economic importance of inefficiencies at oral, descending-price auctions and then compared the expected revenues of the two auction formats.

Finally, we consider a model of fixed costs to bidding and endogenous auction participation, examining Li's (2005) application of a simulated, generalized method-of-moments estimator within a parametric specification.

We devote chapter 5 to an investigation of multi-unit auctions. During the past four decades, economic theorists have systematically investigated simple theoretical models of behavior at auctions in which only one object is sold to buyers demanding at most one object each. In reality, however, many auctions involve the sale of multiple units of the same object to buyers who may demand several units. Recent research concerning multi-unit auction models suggests that such institutions introduce a host of additional economic issues typically absent in the analysis of single-object auction models.

We begin the chapter by first making the distinction between multi-object and multi-unit auctions and then introducing Weber's (1983) classification of multi-unit auctions. Subsequently, we introduce models of singleton demand and then describe two models of multi-unit demand, nonrandom and random demand. For completed research, we then describe identification and estimation strategies. Because multi-unit demand and supply models are topics of current research, our discussion in this chapter is incomplete.

In our last chapter, chapter 6, we discuss briefly directions for future research. We first describe some research that is currently either under revision or under way, and then speculate on a few fruitful directions in which researchers might go. Finally, we summarize the book, briefly.

We have written a number of technical appendixes to this book. We encourage the reader to master the material in them *before* attempting the next four chapters of the book. In these appendixes are included a review of some basic probability theory concerning distributions of transformations of random variables and, particularly, order statistics. We have also presented a brief review of first-order asymptotic methods as well as simulation methods and the bootstrap; the application of these methods to the evaluation of different estimation strategies is also described. The implementation of different estimation strategies is motivated by descriptions of some elementary tools from numerical analysis. Because using numerical methods requires their implementation in some sort of programming environment, our final appendix is a primer concerning MATLAB.

In an effort to provide readers with instruments to gauge their understanding of this material, at the end of this chapter, we have presented several practice problems. Other practice problems, which build on the material covered in the problems at the end of this chapter, are included at the end of each of the next four chapters. A reader who has successfully completed these practice problems will be able to analyze data from an actual auction and then derive policy conclusions from this research. We hope that readers will make this effort and thus enter the exciting field of SEAD.

1.4 Practice Problems

The problems at the end of this chapter are designed to give you some practice with the basics of probability theory as well as statistical estimation and inference that are presumed in the remainder of the book. By implementing the estimation strategies in the programming language MATLAB (or any other programming language for that matter), you will also gain some practice in the elementary numerical methods needed later in the book.

1. Consider a discrete random variable N having probability mass function

$$f_N(n;\theta^0) = \frac{-(\theta^0)^n}{n\log(1-\theta^0)} \qquad n = 1, 2, \ldots,\ 0 < \theta^0 < 1$$

which is often referred to as the *logarithmic series* distribution for reasons that will become clear later in the problem.

   a) Prove that
   $$\sum_{n=1}^{\infty} f_N(n;\theta^0) = 1.$$
   (Hint: consider the Maclaurin-series expansion of $\log(1+x)$ and substitute in $x = -\theta^0$.)

   b) Find the expected value of N, $\mathcal{E}(N)$. (Hint: $\sum_{n=1}^{\infty} \rho^n = \frac{\rho}{1-\rho}$.)

   c) Find the variance of N, $\mathcal{V}(N)$. (Hint: remember that the derivative of a sum is the sum of the derivatives of each of the sum's parts.)

   d) Define the method-of-moments estimator $\hat{\theta}_{\text{MM}}$ of θ^0.

N	1	2	3	4	5	6	7	8	9+
Frequency	700	205	50	26	10	6	1	1	1

Table 1.1
Observed Frequency Distribution of N

e) Show that the condition that defines $\hat{\theta}_{\text{MM}}$ has a unique solution. (Hint: draw a graph.)

f) Set up the recursion you would use in order to employ Newton's method to solve for $\hat{\theta}_{\text{MM}}$.

g) Define the maximum-likelihood estimator $\hat{\theta}_{\text{ML}}$ of θ^0.

h) Demonstrate that $\hat{\theta}_{\text{MM}}$ and $\hat{\theta}_{\text{ML}}$ are consistent estimators of θ^0.

i) Find an approximation to the variance of $\hat{\theta}_{\text{MM}}$ and $\hat{\theta}_{\text{ML}}$.

j) Characterize the asymptotic distribution of $\hat{\theta}_{\text{MM}}$ and $\hat{\theta}_{\text{ML}}$ and explain your reasoning.

After considerable effort, a researcher has obtained a random sample of one thousand measurements on N. These data are summarized in Table 1.1.

k) Write a MATLAB program to implement Newton's method and then calculate the maximum-likelihood estimate of θ^0 using the above data.

l) At size 0.05, test the following hypothesis:

$$\text{H}_0 : \theta^0 = 0.50$$
$$\text{H}_1 : \theta^0 \neq 0.50.$$

m) At size 0.10, test the following hypothesis:

$$\text{H}_0 : \log \theta^0 = -0.70$$
$$\text{H}_1 : \log \theta^0 \neq -0.70.$$

n) At size 0.05, ignoring the fact that $\hat{\theta}_{\text{ML}}$ is estimated and that no observed counts exist above nine, use Fisher's χ^2, *goodness-of-fit* test to decide whether the empirical frequency is consistent with the logarithmic series distribution.

2. Consider a random sample $\{V_t\}_{t=1}^T$ from the log-normal distribution, having probability density function

$$f_V(v;\theta_1^0,\theta_2^0) = \frac{1}{v}\frac{1}{\sqrt{2\pi\theta_2^0}}\exp\left[\frac{-(\log v-\theta_1^0)^2}{2\theta_2^0}\right]$$
$$v>0,\ \theta_2^0>0,\ \text{and} -\infty<\theta_1^0<\infty.$$

Note that the k^{th} raw moment of V is

$$\mathcal{E}(V^k) = \exp\left(k\theta_1^0 + \frac{k^2}{2}\theta_2^0\right) \quad k=1,2,\ldots.$$

a) Write down the likelihood function, the logarithm of the likelihood function, and the score vector for this sample. Solve for the maximum-likelihood estimators $\hat{\theta}_1^{\text{ML}}$ and $\hat{\theta}_2^{\text{ML}}$ of θ_1^0 and θ_2^0.
b) Calculate the expectations of $\hat{\theta}_1^{\text{ML}}$ and $\hat{\theta}_2^{\text{ML}}$. Are the MLEs unbiased estimators? Calculate the variance and the small-sample, exact distribution of $\hat{\theta}_1^{\text{ML}}$.
c) Derive the method-of-moments estimators $\hat{\theta}_1^{\text{MM}}$ and $\hat{\theta}_2^{\text{MM}}$. Are MMEs unbiased estimators of θ_1^0 and θ_2^0? Explain your answer.
d) Prove that the MMEs are consistent estimators of θ_1^0 and θ_2^0.
e) For simplicity, assume that θ_2^0 is known to equal one. Using the delta method, find the asymptotic distribution of $\hat{\theta}_1^{\text{MM}}$. Compare the asymptotic variance of this estimator with the exact variance of the MLE. Which estimator is more efficient? Why?

3. Suppose that, in the model of practice problem 1, θ^0 depends on a $(K\times 1)$ vector of covariates $\boldsymbol{z}$. Assume further that the unknown $\theta(\boldsymbol{z})$ can be modeled as a logistic function, so

$$\theta(\boldsymbol{z}) = \frac{\exp(\boldsymbol{z}^\top\boldsymbol{\gamma})}{[1+\exp(\boldsymbol{z}^\top\boldsymbol{\gamma})]}$$

where the vector of unknown parameters $\boldsymbol{\gamma}$, or $(\gamma_0,\gamma_1,\ldots,\gamma_{K-1})^\top$, is conformable to $\boldsymbol{z}^\top$.

a) For a sample $\{(\boldsymbol{z}_t,n_t)\}_{t=1}^T$, write down the likelihood function, the logarithm of the likelihood function $\mathcal{L}(\boldsymbol{\gamma})$, the score vector $\mathbf{g}(\boldsymbol{\gamma})$, and the Hessian matrix $\mathbf{H}(\boldsymbol{\gamma})$
b) Set up the recursion you would need in order to solve for the maximum-likelihood estimate $\hat{\boldsymbol{\gamma}}$.

c) In the file `logser.dat`, which is located on the CD accompanying this book, you will find five columns of numbers. In the first is recorded an identification number, which ranges from 1 to 1000, while in the second is recorded the dependent variable n_t. In the next three are recorded the covariates $z_{1,t}$, $z_{2,t}$, and $z_{3,t}$. The entire file has $1,000$ rows, so $1,000$ observations. Write a MATLAB program to calculate the maximum-likelihood estimate of $\boldsymbol{\gamma}$ when a constant is present in $\boldsymbol{z}^\top \boldsymbol{\gamma}$.

d) Using the likelihood-ratio test, decide whether the following hypothesis can be rejected at size 0.05:

$$\begin{aligned}
&\mathrm{H}_0 : \gamma_1^0 = \gamma_2^0 = \gamma_3^0 = 0 \\
&\mathrm{H}_1 : \text{not } \mathrm{H}_0.
\end{aligned}$$

4. To get some practice implementing the bootstrap, complete the following:

a) In MATLAB, generate 1,000 samples of size twenty-five for normal pseudo-random variables having mean zero and variance one. For each sample, calculate the sample median and then simulate the nonparametric bootstrap standard error of the sample median using 100 bootstrap samples. Using this information, gauge the accuracy of the asymptotic formula for the variance of the sample median, which is

$$\frac{\pi \sigma^2}{2T}$$

where σ^2 is the variance (one in this case), T is the sample size (twenty-five in this case), and π can be approximated by 3.14159.

b) In MATLAB, generate 1,000 samples of size twenty-five for uniform pseudo-random numbers. Using the property that cumulative distribution function of a continuous random variable is distributed uniformly on the interval $[0, 1]$, generate pseudo-random variables from the exponential distribution having hazard rate one. Using the bootstrap with 100 samples, evaluate the asymptotic formula for the standard error of the sample lower quartile $\hat{\xi}$ as an estimator of the population lower quartile ξ^0 when the asymptotic distribution of the lower quartile is

$$\sqrt{T}(\hat{\xi} - \xi^0) \xrightarrow{\mathrm{d}} \boldsymbol{\mathcal{N}} \left[0, \frac{3}{16 T f_V^0 \left(\xi^0\right)^2} \right]$$

where $f_V^0(\cdot)$ is the true exponential probability density function.

5. Consider the following function

$$f_V^0(v) = \exp\left[v - v^2 - \sqrt{v} + \sin(v)\right] \quad v \in [0, 3].$$

   a) Plot this function in MATLAB.

   In MATLAB, evaluate the above function at each point from 0 to 3 for a constant step-size 0.1; store the thirty-one ordered pairs.

   b) On the interval $[0, 3]$, estimate the generalized Chebyshev polynomial approximations of the function $f_V^0(v)$ for orders one, four, and seven using MATLAB. Graph these three polynomials superimposing them on a graph of the true function $f_V^0(v)$. Do the approximations improve as the order of the approximating polynomial increases?

   c) Using the thirty-one ordered pairs of numbers from part a) above, estimate, by the method of least squares, the appropriate coefficients for a polynomial of order one, four, and seven and then plot these functions along with the true function $f_V^0(v)$.

6. In a variety of circumstances in econometrics, researchers often need to evaluate an integral of the following form:

$$\Gamma(a, b) = \int_a^b f(u) \ du.$$

   In some circumstances $\Gamma(a, b)$ will have a closed-form solution that can be calculated in a straighforward fashion. For example, suppose that

$$f(u) = \exp(-u) \quad u > 0.$$

   In this case,

$$\Gamma(a, b) = [1 - \exp(-b)] - [1 - \exp(-a)] = [\exp(-a) - \exp(-b)].$$

   In many cases, $\Gamma(a, b)$ does not have a closed-form solution. In these cases, quadrature methods are often used to calculate $\Gamma(a, b)$. Quadrature involves dividing the interval $[a, b]$ up into subintervals, evaluating the area under $f(u)$ for each subinterval, and then adding up the areas to find $\Gamma(a, b)$. In higher dimensions (more than three) quadrature rules can become numerically unreliable and difficult to implement with any precision. In such cases, researchers often use Monte Carlo methods to simulate the integral $\Gamma(a, b)$. Monte Carlo simulation involves sampling from a known distribution on the interval $[a, b]$, for example the uniform, and then taking the average of $f(u)$ evaluated at each random draw. As the number of simulation draws K goes to infinity, this estimator converges to the truth.

In this problem, you will use simple trapezoidal quadrature as well as Monte Carlo methods to evaluate

$$\Gamma(0, 1) = \int_0^1 \exp(-u)\ du$$

which you know equals $[1 - \exp(-1)]$ or 0.6321.

a) Approximate $\Gamma(0, 1)$ by the area of a trapezoid defined by the points $(0, 0)$, $(0, 1)$, $(1, 0)$, and $(1, \exp(-1))$, then calculate the error associated with using this rule.
b) Now divide the interval $[0, 1]$ up into ten subintervals of the same width. Calculate the area for each trapezoid, and then the estimated area for $\Gamma(0, 1)$. What is the estimation error now?
c) Derive a formula for the estimation error as a function of the points a and b as well as the number of subintervals on $[a, b]$, assuming that the trapezoid rule is used and that the function $f(u)$ equals $\exp(-u)$.
d) Now consider making independent and identically distributed draws concerning uniform random variable U on the interval $[0, 1]$. Calculate the expected value of $\exp(-U)$ on the interval $[0, 1]$. Calculate the variance of $\exp(-U)$ on the interval $[0, 1]$.
e) Provide a simulation estimator $G_K(0, 1)$ of the integral $\Gamma(0, 1)$ where K is the number of simulation draws. Find its asymptotic distribution.
f) How large must K be before the root mean-squared error of $G_K(0, 1)$ equals the error in part a)?