AUCTION THEORY

VIJAY KRISHNA

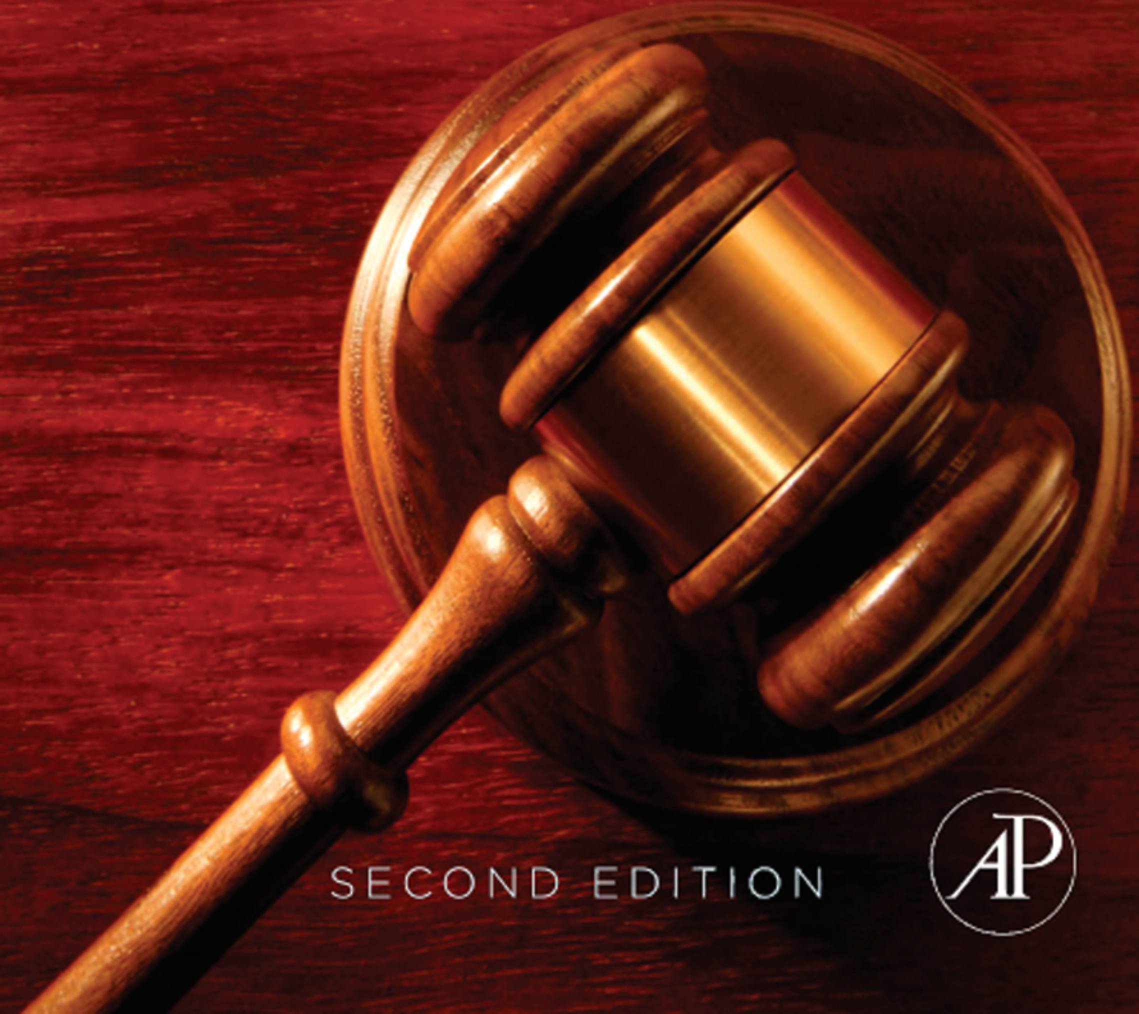

SECOND EDITION

Auction Theory

Second Edition

Vijay Krishna

Department of Economics
Pennsylvania State University
University Park, PA

ELSEVIER

AMSTERDAM • BOSTON • HEIDELBERG • LONDON
NEW YORK • OXFORD • PARIS • SAN DIEGO
SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO
Academic Press is an imprint of Elsevier

Academic Press is an imprint of Elsevier
30 Corporate Drive, Suite 400, Burlington, MA 01803, USA
525 B Street, Suite 1900, San Diego, California 92101-4495, USA
84 Theobald's Road, London WC1X 8RR, UK

This book is printed on acid-free paper. ∞

Library of Congress Cataloging-in-Publication Data
Application Submitted.

British Library Cataloguing in Publication Data
A catalogue record for this book is available from the British Library.

ISBN 13: 978-0-12-374507-1

For information on all Elsevier Academic Press publications
visit our Web site at *www.elsevierdirect.com*

Typeset by: diacriTech, India

Printed in the United States of America
09 10 9 8 7 6 5 4 3 2 1

Contents

Preface . . . . **ix**

1 Introduction **1**
1.1 Some Common Auction Forms 2
1.2 Valuations 2
1.3 Equivalent Auctions 4
1.4 Revenue versus Efficiency 5
1.5 What Is an Auction? 6
1.6 Outline of Part I 6

I Single-Object Auctions 9

2 Private Value Auctions: A First Look **11**
2.1 The Symmetric Model 11
2.2 Second-Price Auctions 12
2.3 First-Price Auctions 13
2.4 Revenue Comparison 17
2.5 Reserve Prices 21

3 The Revenue Equivalence Principle **27**
3.1 Main Result 27
3.2 Some Applications of the Revenue Equivalence Principle 29
3.2.1 Unusual Auctions 29
3.2.2 Uncertain Number of Bidders 32

4 Qualifications and Extensions **37**
4.1 Risk-Averse Bidders 38
4.2 Budget Constraints 42
4.2.1 Second-Price Auctions 42
4.2.2 First-Price Auctions 44
4.2.3 Revenue Comparison 45
4.3 Asymmetries among Bidders 46
4.3.1 Asymmetric First-Price Auctions with Two Bidders 46
4.3.2 Revenue Comparison 51
4.3.3 Efficiency Comparison 53
4.4 Resale and Efficiency 54

5 Mechanism Design 61
5.1 Mechanisms 61
5.1.1 The Revelation Principle 62
5.1.2 Incentive Compatibility 63
5.1.3 Individual Rationality 66
5.2 Optimal Mechanisms 67
5.2.1 Setup 67
5.2.2 Solution 68
5.2.3 Discussion and Interpretation 71
5.2.4 Auctions versus Mechanisms 74
5.3 Efficient Mechanisms 75
5.3.1 The VCG Mechanism 75
5.3.2 Budget Balance 77
5.3.3 An Application to Bilateral Trade 79

6 Auctions with Interdependent Values 85
6.1 The Symmetric Model 88
6.2 Second-Price Auctions 89
6.3 English Auctions 91
6.4 First-Price Auctions 94
6.5 Revenue Comparisons 97
6.5.1 English versus Second-Price Auctions 97
6.5.2 Second-Price versus First-Price Auctions 98
6.6 Efficiency 101

7 The Revenue Ranking (“Linkage”) Principle 105
7.1 The Main Result 105
7.2 Public Information 108
7.3 An Alternative Linkage Principle 110

8 Asymmetries and Other Complications 113
8.1 Failures of the Linkage Principle 113
8.2 Asymmetric Equilibria in Symmetric Second-Price Auctions 118
8.3 Asymmetrically Informed Bidders 120
8.4 Reserve Prices and Entry Fees 122

9 Efficiency and the English Auction 129
9.1 The Single Crossing Condition 130
9.2 Two-Bidder Auctions 131
9.3 The Average Crossing Condition 134
9.4 Three or More Bidders 135
9.5 Proof of Proposition 9.2 136
9.6 Miscellany 141

10 Mechanism Design with Interdependent Values 145
10.1 Efficient Mechanisms 146
10.2 Optimal Mechanisms 150

11 Bidding Rings 157
11.1 Collusion in Second-Price Auctions 158
11.1.1 Efficient Collusion 160
11.1.2 Reserve Prices in the Face of Collusion 163
11.2 Collusion in First-Price Auctions 166

II Multiple-Object Auctions 171

12 An Introduction to Multiple-Object Auctions 173
12.1 Sealed-Bid Auctions for Selling Identical Units 174
12.1.1 Discriminatory Auctions 176
12.1.2 Uniform-Price Auctions 177
12.1.3 Vickrey Auctions 178
12.2 Some Open Auctions 179
12.2.1 Dutch Auctions 179
12.2.2 English Auctions 180
12.2.3 Ausubel Auctions 180

13 Equilibrium and Efficiency with Private Values 185
13.1 The Basic Model 185
13.2 Vickrey Auctions 187
13.3 Efficiency in Multiunit Auctions 189
13.4 Uniform-Price Auctions 190
13.4.1 Demand Reduction 192
13.4.2 Single-Unit Demand 196
13.5 Discriminatory Auctions 196
13.5.1 Structure of Equilibria 197
13.5.2 Single-Unit Demand 200

14 Some Revenue Considerations 203
14.1 Revenue Equivalence in Multiunit Auctions 204
14.2 Revenue Equivalence with Multiunit Demand: An Example .. 206

15 Sequential Sales 213
15.1 Sequential First-Price Auctions 213
15.1.1 Two Units 214
15.1.2 More than Two Units 217
15.1.3 Equilibrium Bids and Prices 219

Chapter | One

Introduction

In A.D. 193, having killed the Emperor Pertinax, in a bold move the Prætorian Guard proceeded to sell off the entire Roman Empire by means of an auction. The winning bid was a promise of 25,000 sesterces per man to the Guard. The winner, Didius Julianus, was duly declared emperor but lasted for only two months before suffering from what is perhaps the earliest and most extreme instance of the "winner's curse": He was beheaded.

Auctions have been used since antiquity for the sale of a variety of objects. Herodotus reports that auctions were used in Babylon as early as 500 B.C. Today, both the range and the value of objects sold by auction have grown to staggering proportions. Art objects and antiques have always been sold at the fall of the auctioneer's hammer. But now numerous kinds of commodities, ranging from tobacco, fish, and fresh flowers to scrap metal and gold bullion, are sold by means of auctions. Bond issues by public utilities are usually auctioned off to investment banking syndicates. Long-term securities are sold in weekly auctions conducted by the U.S. Treasury to finance the borrowing needs of the government. Perhaps the most important use of auctions has been to facilitate the transfer of assets from public to private hands—a worldwide phenomenon in the past two decades. These have included the sale of industrial enterprises in Eastern Europe and the former Soviet Union, and transportation systems in Britain and Scandinavia. Traditionally, the rights to use natural resources from public property—such as timber rights and offshore oil leases—have been sold by means of auctions. In the modern era, auctions of rights to use the electromagnetic spectrum for communication are also a worldwide phenomenon. Finally, there has been a tremendous growth in both the number of Internet auction websites, where individuals can put up items for sale under common auction rules, and the value of goods sold there.

The process of procurement via competitive bidding is nothing but an auction, except that in this case the bidders compete for the right to sell their products or services. Billions of dollars of government purchases are almost exclusively made in this way, and the practice is widespread, if not endemic, in business.

In what follows, an auction will be understood to include the process of procurement via competitive bidding. Of course, in this case it is the person bidding lowest who wins the contract.

Why are auctions and competitive bidding so prevalent? Are there situations to which an auction is particularly suited as a selling mechanism as opposed to, say, a fixed, posted price? From the point of view of the bidders, what are good bidding strategies? From the point of view of the sellers, are particular forms of auctions likely to bring greater revenues than others? These and other questions form the subject matter of this book.

1.1 SOME COMMON AUCTION FORMS

The open ascending price or *English* auction is the oldest and perhaps most prevalent auction form. The word *auction* itself is derived from the Latin *augere*, which means "to increase" (or "augment"), via the participle *auctus* ("increasing"). In one variant of the English auction, the sale is conducted by an auctioneer who begins by calling out a low price and raises it, typically in small increments, as long as there are at least two interested bidders. The auction stops when there is only one interested bidder. One way to formally model the underlying game is to postulate that the price rises continuously and each bidder indicates an interest in purchasing at the current price in a manner apparent to all by, say, raising a hand. Once a bidder finds the price to be too high, he signals that he is no longer interested by lowering his hand. The auction ends when only a single bidder is still interested. This bidder wins the object and pays the auctioneer an amount equal to the price at which the second-last bidder dropped out.

The *Dutch* auction is the open descending price counterpart of the English auction. It is not commonly used in practice but is of some conceptual interest. Here, the auctioneer begins by calling out a price high enough so that presumably no bidder is interested in buying the object at that price. This price is gradually lowered until some bidder indicates her interest. The object is then sold to this bidder at the given price.

The sealed-bid *first-price* auction is another common form. Its workings are rather straightforward: Bidders submit bids in sealed envelopes; the person submitting the highest bid wins the object and pays what he bid.

Finally, there is the sealed-bid *second-price* auction. As its name suggests, once again bidders submit bids in sealed envelopes; the person submitting the highest bid wins the object but pays not what he bid but the second-highest bid.

1.2 VALUATIONS

Auctions are used precisely because the seller is unsure about the values that bidders attach to the object being sold—the maximum amount each bidder is willing to pay. If the seller knew the values precisely, he could just offer the object to the bidder with the highest value at or just below what this bidder is willing to pay. The uncertainty regarding values facing both sellers and buyers is an inherent feature of auctions.

If each bidder knows the value of the object to himself at the time of bidding, the situation is called one of privately known values or *private values*. Implicit in this situation is that no bidder knows with certainty the values attached by *other* bidders and knowledge of other bidders' values would not affect how much the object is worth to a particular bidder. The assumption of private values is most plausible when the value of the object to a bidder is derived from its consumption or use alone. For instance, if bidders assign different values to a painting, a stamp, or a piece of furniture only on the basis of how much utility they would derive from possessing it, perhaps viewing it purely as a consumption good, then the private values assumption is reasonable. On the other hand, if bidders assign values on the basis of how much the object will fetch in the resale market, viewing it as an investment, then the private values assumption is not a good one.

In many situations, how much the object is worth is unknown at the time of the auction to the bidder himself. He may have only an estimate of some sort or some privately known signal—such as an expert's estimate or a test result—that is correlated with the true value. Indeed, other bidders may possess information—say, additional estimates or test results—that if known, would affect the value that a particular bidder attaches to the object. Thus, values are unknown at the time of the auction and may be affected by information available to other bidders. Such a specification is called one of *interdependent values* and is particularly suited for situations in which the object being sold is an asset that can possibly be resold after the auction. A special case of this is a situation in which the value, though unknown at the time of bidding, is the same for all bidders—a situation described as being one of a pure *common value*.[1] A common value model is most appropriate when the value of the object being auctioned is derived from a market price that is unknown at the time of the auction. An archetypal example is the sale of a tract of land with an unknown amount of oil underground. Bidders may have different estimates of the amount of oil, perhaps based on privately conducted tests, but the final value of the land is derived from the future sales of the oil, so this value is, to a first approximation, the same for all bidders.

Note that the term *interdependence* refers only to the structure of values and how these are affected by information held by other bidders. It does not refer to any statistical properties of this information—that is, how the signals observed by the bidders are distributed. Thus, we could have a situation in which values are interdependent so a particular bidder's value depends on a signal observed by another bidder, but at the same time, the signals themselves are statistically independent. Similarly, we could have a situation in which the values are not interdependent so a particular bidder's value depends only on his own signal, but the signals themselves are correlated.

---

[1] Sometimes the term *common values* is itself used to label what we have called *interdependent values*. We use the latter term because it more accurately describes the situation.

1.3 EQUIVALENT AUCTIONS

Four auction formats have been outlined here. Two were open auctions—the English and the Dutch—and two were sealed-bid auctions—the first- and second-price formats. These seem very different institutions, and certainly, they differ in the way that they are implemented in the real world. Open auctions require that the bidders collect in the same place, whereas sealed bids may be submitted by mail, so a bidder may observe the behavior of other bidders in one format and not in another. For rational decision makers, however, some of these differences are superficial.

First, observe that the Dutch open descending price auction is strategically equivalent to the first-price sealed-bid auction.[2] In a first-price sealed-bid auction, a bidder's strategy maps his private information into a bid. Although the Dutch auction is conducted in the open, it offers no useful information to bidders. The only information that is available is that some bidder has agreed to buy at the current price, but that causes the auction to end. Bidding a certain amount in a first-price sealed-bid auction is equivalent to offering to buy at that amount in a Dutch auction, provided the item is still available. For every strategy in a first-price auction there is an equivalent strategy in the Dutch auction and vice versa.

Second, when values are private, the English open ascending auction is also equivalent to the second-price sealed-bid auction, but in a weaker sense than noted earlier. The English auction offers information about when other bidders drop out, and by observing this, it may be possible to infer something about their privately known information. With private values, however, this information is of no use. In an English auction, it clearly cannot be optimal to stay in after the price exceeds the value—which can only cause a loss—or to drop out before the price reaches the value—thus forgoing potential gains. Likewise, in a second-price auction it is best to bid the value (this is discussed in more detail later). Thus, with private values, the optimal strategy in both is to bid up to or stay in until the value is reached.

This equivalence between the English and second-price auctions is *weak* in two senses. First, the two auctions are not strategically equivalent. Second, and more important, the optimal strategies in the two are the same only if values are private. With interdependent values, the information available to others is relevant to a particular bidder's evaluation of the worth of the object. Seeing some other bidder drop out early may bring bad news that may cause a bidder to reduce his own estimate of the object's value. Thus, if values are interdependent, the two auctions need not be equivalent from the perspective of the bidders. Figure 1.1 depicts the equivalences between the open and sealed-bid formats introduced here.

---

[2]Two games are strategically equivalent if they have the same normal form except for duplicate strategies. Roughly this means that for every strategy in one game, a player has a strategy in the other game, which results in the same outcomes.

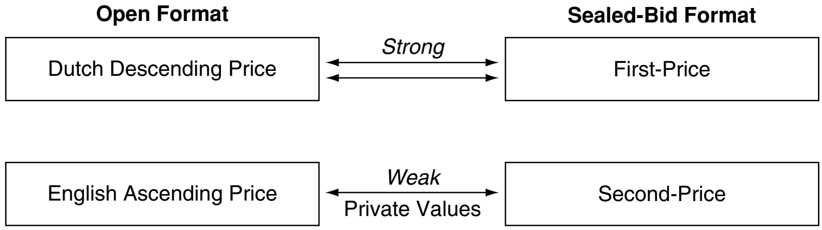

FIGURE 1.1 Equivalence of open and sealed-bid formats.

1.4 REVENUE VERSUS EFFICIENCY

The main questions that guide auction theory involve a comparison of the performance of different auction formats as economic institutions. These are evaluated on two grounds, and the relevance of one or the other criterion depends on the context. From the perspective of the seller, a natural yardstick in comparing different auction forms is the *revenue*, or the expected selling price, that they fetch. From the perspective of society as a whole, however, *efficiency*—that the object end up in the hands of the person who values it the most *ex post*—may be more important. This is especially true when the auction concerns the sale of a publicly held asset to the private sector, so the seller, in this case a government, may want to choose a format that ensures that the object is allocated efficiently, even if the revenue from some other, inefficient format is higher.

But should efficiency be a criterion at all? Why can we not rely on "the market" to reallocate the object efficiently, even if the auction does not do so? After all, if there are unrealized gains from trade, the person who wins the auction can resell the object to someone who attaches a higher value. We will argue that this argument is suspect for many reasons. First, postauction transactions will typically involve a small number of agents, especially in the context of privatization, and so will result in some bargaining about the resale price. Such bargaining is unlikely to result in efficient outcomes, since it will typically take place under conditions of incomplete information. Second, resale may involve significant transaction costs, so it may not take place even when it should. In Chapter 4 we take up the question of whether resale will lead to efficiency more formally. In short, we find that even in the best circumstances—with no transaction costs or bargaining delays—the answer is no. Resale cannot guarantee efficiency, so a policy maker interested in achieving efficiency would do well to choose the auction format carefully.

Of course, revenue and efficiency are not the only criteria that should guide the choice of an auction format. The common auction forms discussed thus far have the virtue of simplicity—the rules of the auction are transparent—and this may be an important practical consideration. Another important factor may be the potential for collusion among bidders. As we will see later, auction formats differ in their susceptibility to such collusion.

1.5 WHAT IS AN AUCTION?

A wide variety of selling institutions fall under the rubric of "an auction." There are hybrid Dutch-English auctions in which the price is lowered until there is an interested bidder and then other bidders are allowed to outbid this amount. There are what may be called "deadline" auctions—commonly used by Internet auction sites—in which the person with the highest standing bid before a fixed stopping time—say, noon on Sunday—is declared the winner. There are "candle" auctions, with a random stopping time, in which the person with the highest bid standing before the wick of a candle burns out wins. One may conceive of a third-price auction or an auction in which the winner pays the average of all the other bids. The range of possibilities is rather wide and even more so when sales of multiple objects are considered. Without adopting a rigid view as to what may be called an auction and what may not, we seek to identify some important features that such institutions have in common.

A common aspect of auction-like institutions is that they elicit information, in the form of bids, from potential buyers regarding their willingness to pay, and the outcome—that is, who wins what and pays how much—is determined solely on the basis of the received information. An implication of this is that auctions are *universal* in the sense that they may be used to sell any good. A valuable piece of art and a secondhand car can both be sold by means of an English auction under the same basic set of rules. Alternatively, both can be sold by means of a first-price sealed-bid auction. The auction form does not depend on any details specific to the item at hand.

A second important aspect of auction-like institutions is that they are *anonymous*. By this we mean that the identities of the bidders play no role in determining who wins the object and who pays how much. So if bidder 1 wins with a bid of b_1 and pays some amount p, then keeping all other bids fixed, if some other bidder—say, bidder 2—were to bid b_1 and bidder 1 were to bid b_2, then bidder 2 would win and pay p also. Every bidder other than 1 and 2—say, bidder 3—is completely unaffected if bidders 1 and 2 exchange their bids in the manner just described.

In later chapters we place auctions in a larger class of institutions, called *mechanisms*. Mechanisms differ from auctions in that they are not necessarily universal or anonymous.

1.6 OUTLINE OF PART I

Part I presents situations where a *single* indivisible object is sold to one of many potential buyers. Chapter 2 introduces the basic theory of auctions with *private values*, beginning with the case where these are symmetrically and independently distributed. It derives equilibrium strategies in first- and second-price auctions and compares their performance. Chapter 3 concerns the benchmark "revenue equivalence principle," in its simplest form. Chapter 4 is then concerned with amendments to the revenue equivalence principle necessitated by various extensions to the basic model including asymmetries, risk aversion, and budget constraints. Chapter 5 examines the problem of mechanism design with private values, considering both optimal and efficient mechanisms.

Chapter 6 introduces the model of auctions with *interdependent values* and affiliated signals, again deriving equilibrium strategies in the common auction forms. The main goal here is to rank the common auction forms in terms of the expected selling price. Chapter 7 derives the "revenue ranking principle" and explores some of its implications. Chapter 8 again explores some extensions and qualifications to the basic model necessitated by asymmetries among bidders. Chapter 9 considers the problem of allocating efficiently when bidders are asymmetric, focusing on the efficiency properties of the English auction. Chapter 10 studies mechanism design with interdependent values, again considering both optimal and efficient mechanisms.

Finally, Chapter 11 is concerned with collusive behavior among bidders and the formation of bidding cartels. The models here are with private values.

Figure 1.2 shows the organization of Part I, emphasizing the more or less parallel development of the subject matter in the private value and the interdependent value cases.

Part II of the book concerns *multiple*-object auctions. Chapter 12 serves as an introduction to this part.

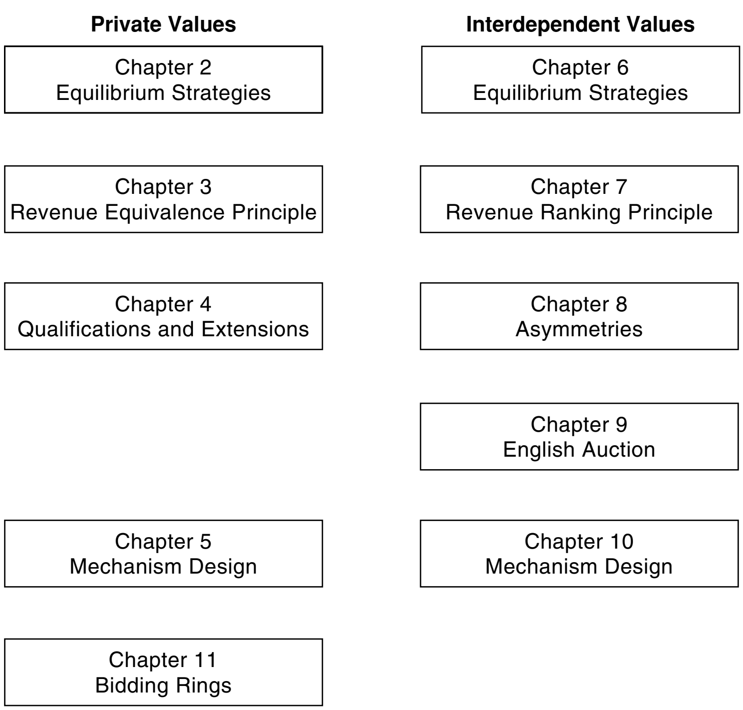

FIGURE 1.2 Outline of Part I.

CHAPTER NOTES

Cassady (1967) provides a panoramic view of real-world auction institutions, past and present, that is both colorful and insightful. Second-price auctions are also referred to as Vickrey auctions. It was commonly believed that the second-price auction was a purely theoretical construct proposed by Vickrey (1961) as a sealed-bid counterpart of the open ascending-price format. Lucking-Reiley (2000) points out, however, that many stamp auctions have been conducted under second-price rules since the nineteenth century. In this context, they originated as a means of allowing bidders who could not be present at the actual, open ascending-price auction, to submit bids by mail.

Many Internet auction websites have adopted what are effectively second-price rules. For instance, at the popular auction site eBay, goods are sold by means of what appears to be an English auction. Bidders can, however, make use of *proxy bidding* wherein they employ a computer program, sometimes called an "elf," to bid on their behalf. The computer program raises rival bids by the minimum increment as long as it is below some limit set by the bidder. It is easy to see that this is effectively a second-price auction in which the amount bid is the same as the limit set by a bidder. Again, see the paper by Lucking-Reiley (2000).

There have been many excellent surveys of auction theory. These vary in both content and emphasis, reflecting, as does this book, the interests of the authors and the state of theory at the time they were written. We mention some of the prominent ones. Milgrom (1985) gives a cogent account of the theory of symmetric single-object auctions and shows how the theory may be extended to situations in which there are multiple objects but each bidder wants at most only one. McAfee and McMillan (1987a) also concentrate on the symmetric single object case but emphasize many extensions and applications of the theory. Milgrom (1987) attempts to answer the question of when auctions are appropriate and why they are so prevalent. He places auctions in the larger context of general institutions of economic exchange and evaluates their performance in different environments. The survey by Wilson (1992)—again largely concerning single-object auctions—offers a wide range of examples in which equilibrium bidding strategies can be computed in closed form. Technical aspects of the symmetric private values model are carefully treated by Matthews (1995). Klemperer (2003) emphasizes that many aspects of auction theory have interesting applications to other branches of economic theory.

There is now a substantial and rapidly growing literature concerning empirical work on auctions and the development of associated econometric tools. A detailed discussion would take us too far afield, so we only mention a representative sample of the work. The papers by Hendricks, Porter, and Wilson (1994), Hendricks and Paarsch (1995), and Laffont, Ossard, and Vuong (1995) serve as useful introductions to the area.

Auctions have also been the subject of a now large body of work in experimental economics. Kagel (1995) has written a thoughtful survey of the area.

PART I

Single-Object Auctions

Chapter | Two

Private Value Auctions: A First Look

We begin the formal analysis by considering equilibrium bidding behavior in the four common auction forms in an environment with independently and identically distributed private values. In the previous chapter we argued that the open descending price (or Dutch) auction is strategically equivalent to the first-price sealed-bid auction. When values are private, the open ascending price (or English) auction is also equivalent to the second-price sealed-bid auction, albeit in a weaker sense. Thus, for our purposes, it is sufficient to consider the two sealed-bid auctions.

This chapter introduces the basic methodology of auction theory. We postulate an informational environment consisting of (1) a valuation structure for the bidders—in this case, that of private values—and (2) a distribution of information available to the bidders—in this case, it is independently and identically distributed. We consider different auction formats—in this case, first- and second-price sealed-bid auctions. Each auction format now determines a game of incomplete information among the bidders and, keeping the informational environment fixed, we determine a Bayesian-Nash equilibrium for each resulting game. When there are many equilibria, we usually select one on some basis—dominance, perfection, or symmetry—but make sure that the criterion is applied uniformly to all formats. The relative performance of the auction formats on grounds of revenue or efficiency is then evaluated by comparing the equilibrium outcomes in one format versus another.

2.1 THE SYMMETRIC MODEL

There is a single object for sale, and N potential buyers are bidding for the object. Bidder i assigns a value of X_i to the object—the maximum amount a bidder is willing to pay for the object. Each X_i is independently and identically distributed on some interval $[0, \omega]$ according to the increasing distribution function F. It is assumed that F admits a continuous density $f \equiv F'$ and has full support. We allow for the possibility that the support of F is the nonnegative real line $[0, \infty)$ and if that is so, with a slight abuse of notation, write $\omega = \infty$. In any case, it is assumed that $E[X_i] < \infty$.

Bidder i knows the realization x_i of X_i and only that other bidders' values are independently distributed according to F. Bidders are risk neutral; they seek to maximize their expected profits. All components of the model other than the realized values are assumed to be commonly known to all bidders. In particular, the distribution F is common knowledge, as is the number of bidders.

Finally, it is also assumed that bidders are not subject to any liquidity or budget constraints. Each bidder i has sufficient resources so if necessary, he or she can pay the seller up to his or her value x_i. Thus, each bidder is both willing and able to pay up to his or her value.

We emphasize that the distribution of values is the same for all bidders, and we will refer to this situation as one involving *symmetric* bidders.

In this framework, we examine two major auction formats:

I. A first-price sealed-bid auction, where the highest bidder gets the object and pays the amount he bid

II. A second-price sealed-bid auction, where the highest bidder gets the object and pays the second highest bid

Each of these auction formats determines a game among the bidders. A strategy for a bidder is a function $\beta_i : [0, \omega] \to \mathbb{R}_+$, which determines his or her bid for any value. We will typically be interested in comparing the outcomes of a symmetric equilibrium—an equilibrium in which all bidders follow the same strategy—of one auction with a symmetric equilibrium of the other. Given that bidders are symmetric, it is natural to focus attention on symmetric equilibria. We ask the following questions:

What are symmetric equilibrium strategies in a first-price auction (I) and a second-price auction (II)?

From the point of view of the seller, which of the two auction formats yields a higher expected selling price in equilibrium?

2.2 SECOND-PRICE AUCTIONS

Although the first-price auction format is more familiar and even natural, we begin our analysis by considering second-price auctions. The strategic problem confronting bidders in second-price auctions is much simpler than that in first-price auctions, so they constitute a natural starting point. Also recall that in the private values framework, second-price auctions are equivalent to open ascending price (or English) auctions.

In a second-price auction, each bidder submits a sealed bid of b_i, and given these bids, the payoffs are:

$$\Pi_i = \begin{cases} x_i - \max_{j \neq i} b_j & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{if } b_i < \max_{j \neq i} b_j \end{cases}$$

We also assume that if there is a tie, so $b_i = \max_{j \neq i} b_j$, the object goes to each winning bidder with equal probability. Bidding behavior in a second-price auction is straightforward.

Proposition 2.1. *In a second-price sealed-bid auction, it is a weakly dominant strategy to bid according to $\beta^{\mathrm{II}}(x) = x$.*

Proof. Consider bidder 1, say, and suppose that $p_1 = \max_{j \neq 1} b_j$ is the highest competing bid. By bidding x_1, bidder 1 will win if $x_1 > p_1$ and not if $x_1 < p_1$ (if $x_1 = p_1$, bidder 1 is indifferent between winning and losing). Suppose, however, that he bids an amount $z_1 < x_1$. If $x_1 > z_1 \geq p_1$, then he still wins, and his profit is still $x_1 - p_1$. If $p_1 > x_1 > z_1$, he still loses. However, if $x_1 > p_1 > z_1$, then he loses, whereas if he had bid x_1, he would have made a positive profit. Thus, bidding less than x_1 can never increase his profit but in some circumstances may actually decrease it. A similar argument shows that it is not profitable to bid more than x_1. ■

It should be noted that the argument in Proposition 2.1 relied neither on the assumption that bidders' values were independently distributed nor the assumption that they were identically so. Only the assumption of private values is important, and Proposition 2.1 holds as long as this is the case.

With Proposition 2.1 in hand, let us ask how much each bidder expects to pay in equilibrium. Fix a bidder—say, 1—and let the random variable $Y_1 \equiv Y_1^{(N-1)}$ denote the highest value among the $N-1$ remaining bidders. In other words, Y_1 is the highest-order statistic of $X_2, X_3, \ldots, X_N$ (see Appendix C). Let G denote the distribution function of Y_1. Clearly, for all y, $G(y) = F(y)^{N-1}$. In a second-price auction, the expected payment by a bidder with value x can be written as

$$\begin{aligned} m^{\mathrm{II}}(x) &= \text{Prob[Win]} \times E[\text{2nd highest bid} \mid x \text{ is the highest bid}] \\ &= \text{Prob[Win]} \times E[\text{2nd highest value} \mid x \text{ is the highest value}] \\ &= G(x) \times E[Y_1 \mid Y_1 < x] \end{aligned} \tag{2.1}$$

2.3 FIRST-PRICE AUCTIONS

In a first-price auction, each bidder submits a sealed bid of b_i, and given these bids, the payoffs are

$$\Pi_i = \begin{cases} x_i - b_i & \text{if } b_i > \max_{j \neq i} b_j \\ 0 & \text{if } b_i < \max_{j \neq i} b_j \end{cases}$$

As before, if there is more than one bidder with the highest bid, the object goes to each such bidder with equal probability.

In a first-price auction, equilibrium behavior is more complicated than in a second-price auction. Clearly, no bidder would bid an amount equal to his or her value, since this would only guarantee a payoff of 0. Fixing the bidding behavior of others, at any bid that will neither win for sure nor lose for sure, the bidder faces a simple trade-off. An increase in the bid will increase the probability of winning while, at the same time reducing the gains from winning. To get some idea about how these effects balance off, we begin with a heuristic derivation of symmetric equilibrium strategies.

Suppose that bidders $j \neq 1$ follow the symmetric, increasing, and differentiable equilibrium strategy $\beta^{\mathrm{I}} \equiv \beta$. Suppose bidder 1 receives a signal, $X_1 = x$, and bids b. We wish to determine the optimal b.

First, notice that it can never be optimal to choose a bid $b > \beta(\omega)$, since in that case, bidder 1 would win for sure and could do better by reducing his bid slightly, so he still wins for sure but pays less. So we need only consider bids $b \leq \beta(\omega)$. Second, a bidder with value 0 would never submit a positive bid, since he would make a loss if he were to win the auction. Thus, we must have $\beta(0) = 0$.

Bidder 1 wins the auction whenever he submits the highest bid—that is, whenever $\max_{i\neq 1} \beta(X_i) < b$. Since β is increasing, $\max_{i\neq 1} \beta(X_i) = \beta(\max_{i\neq 1} X_i) = \beta(Y_1)$, where, as before, $Y_1 \equiv Y_1^{(N-1)}$, the highest of $N-1$ values. Bidder 1 wins whenever $\beta(Y_1) < b$ or equivalently, whenever $Y_1 < \beta^{-1}(b)$. His expected payoff is therefore

$$G\left(\beta^{-1}(b)\right) \times (x - b)$$

where, again, G is the distribution of Y_1. Maximizing this with respect to b yields the first-order condition:

$$\frac{g\left(\beta^{-1}(b)\right)}{\beta'\left(\beta^{-1}(b)\right)}(x-b) - G(\beta^{-1}(b)) = 0 \tag{2.2}$$

where $g = G'$ is the density of Y_1.

At a symmetric equilibrium, $b = \beta(x)$, and thus (2.2) yields the differential equation

$$G(x)\beta'(x) + g(x)\beta(x) = xg(x) \tag{2.3}$$

or equivalently,

$$\frac{d}{dx}(G(x)\beta(x)) = xg(x)$$

and since $\beta(0) = 0$, we have

$$\begin{aligned}\beta(x) &= \frac{1}{G(x)} \int_0^x y g(y)\, dy \\ &= E[Y_1 \mid Y_1 < x]\end{aligned}$$

The derivation of β is only heuristic because (2.3) is merely a necessary condition: We have not formally established that if the other $N-1$ bidders follow β, then it is indeed optimal for a bidder with value x to bid $\beta(x)$. The next proposition verifies that this is indeed correct.

Proposition 2.2. *Symmetric equilibrium strategies in a first-price auction are given by*

$$\beta^{\mathrm{I}}(x) = E[Y_1 \mid Y_1 < x] \tag{2.4}$$

where Y_1 is the highest of $N-1$ independently drawn values.

Proof. Suppose that all but bidder 1 follow the strategy $\beta^{\mathrm{I}} \equiv \beta$ given in (2.4). We will argue that in that case it is optimal for bidder 1 to follow β also. First, notice that β is an increasing and continuous function. Thus, in equilibrium the bidder with the highest value submits the highest bid and wins the auction. It is not optimal for bidder 1 to bid a $b > \beta(\omega)$. The expected payoff of bidder 1 with value x if he bids an amount $b \leq \beta(\omega)$ is calculated as follows. Denote by $z = \beta^{-1}(b)$ the value for which b is the equilibrium bid—that is, $\beta(z) = b$. Then we can write bidder 1's expected payoff from bidding $\beta(z)$ when his value is x as follows:

$$\begin{aligned}\Pi(b,x) &= G(z)[x - \beta(z)] \\ &= G(z)x - G(z)E[Y_1 \mid Y_1 < z] \\ &= G(z)x - \int_0^z y g(y)\, dy \\ &= G(z)x - G(z)z + \int_0^z G(y)\, dy \\ &= G(z)(x - z) + \int_0^z G(y)\, dy\end{aligned}$$

where the fourth equality is obtained as a result of integration by parts. (Alternatively, see formula (A.2) in Appendix A.)

We thus obtain that

$$\Pi(\beta(x),x) - \Pi(\beta(z),x) = G(z)(z - x) - \int_x^z G(y)\, dy \geq 0$$

regardless of whether $z \geq x$ or $z \leq x$.

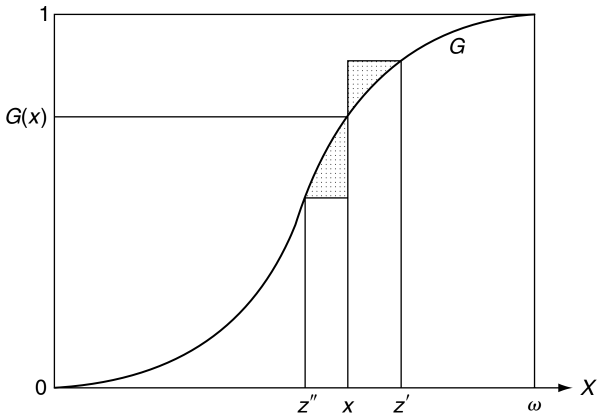

FIGURE 2.1 Losses from over- and underbidding in a first-price auction.

(The preceding argument shows that bidding an amount $\beta(z') > \beta(x)$ rather than $\beta(x)$ results in a loss equal to the shaded area to the right in Figure 2.1; similarly, bidding an amount $\beta(z'') < \beta(x)$ results in a loss equal to the area to the left.)

We have thus argued that if all other bidders are following the strategy β, a bidder with a value of x cannot benefit by bidding anything other than $\beta(x)$, and this implies that β is a symmetric equilibrium strategy. ■

The equilibrium bid can be rewritten as

$$\beta^{\mathrm{I}}(x) = x - \int_0^x \frac{G(y)}{G(x)}\, dy$$

by using (A.2) in Appendix A again. This shows that the bid is, naturally, less than the value x. Since

$$\frac{G(y)}{G(x)} = \left[\frac{F(y)}{F(x)}\right]^{N-1}$$

the degree of "shading" (the amount by which the bid is less than the value) depends on the number of competing bidders and as N increases, approaches 0. Thus, for fixed F, as the number of bidders increases, the equilibrium bid $\beta^{\mathrm{I}}(x)$ approaches x.

It is instructive to derive the equilibrium strategies explicitly in a few examples.

Example 2.1. *Values are uniformly distributed on* $[0, 1]$.

If $F(x) = x$, then $G(x) = x^{N-1}$ and

$$\beta^{\mathrm{I}}(x) = \frac{N-1}{N} x$$

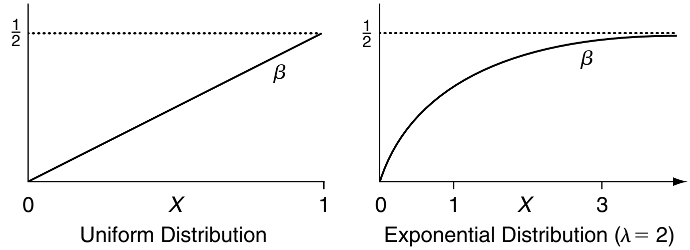

FIGURE 2.2 Equilibria of two-bidder symmetric first-price auctions.

In this case, the equilibrium strategy calls upon a bidder to bid a constant fraction of his value. For the case of two bidders, the equilibrium bidding strategy is depicted in the left-hand panel of Figure 2.2. ▲

Example 2.2. *Values are exponentially distributed on* $[0,\infty)$*, and there are only two bidders.*

If $F(x)=1-\exp(-\lambda x)$, for some $\lambda>0$, and $N=2$, then

$$\begin{aligned}\beta^{\mathrm{I}}(x) &= x-\int_0^x \frac{F(y)}{F(x)}\,dy\\ &= \frac{1}{\lambda}-\frac{x\exp(-\lambda x)}{1-\exp(-\lambda x)}\end{aligned}$$

As a particular instance, consider the case where $\lambda=2$ so that $E[X]=\frac{1}{2}$. The equilibrium bidding strategy in this case is depicted in the right-hand panel of Figure 2.2. The figure highlights the fact that with the exponentially distributed values, even a bidder with a very high value—say, \$1 million—will not bid more than 50 cents! This seems counterintuitive at first—the bidder is facing the risk of a big loss by not bidding higher—but is explained by the fact that the probability that the bidder with a high value will lose in equilibrium is infinitesimal. Indeed, for a bidder with a value of \$1 million, it is smaller than $10^{-400000}$. This fact, together with the assumption that bidders are risk neutral, implies that bidders with high values are willing to bid very small amounts. Formally, the fact that no bidder bids more than $\frac{1}{2}$ is a consequence of the property that for all x,

$$\beta^{\mathrm{I}}(x)=E\left[Y_1 \mid Y_1<x\right]\leq E\left[Y_1\right]$$

and when there are only two bidders, the latter is the same as $E[X]$. ▲

2.4 REVENUE COMPARISON

Having derived symmetric equilibrium strategies in both the second- and first-price auctions, we can now compare the selling prices—the revenues accruing to the seller—in the two formats.

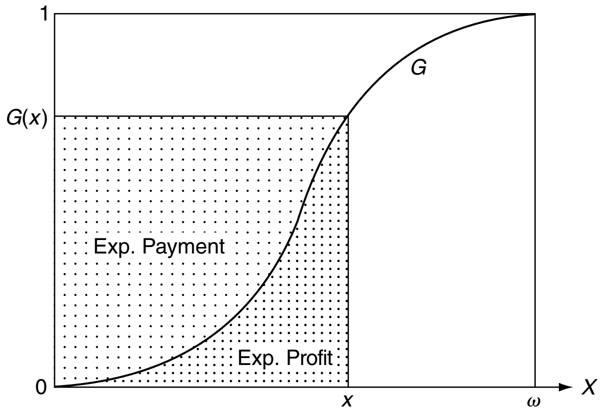

FIGURE 2.3 Payments and profits in first- and second-price auctions.

In a first-price auction, the winner pays what he or she bid, and thus the expected payment by a bidder with value x is

$$m^{\mathrm{I}}(x) = \text{Prob[Win]} \times \text{Amount bid} = G(x) \times E[Y_1 \mid Y_1 < x] \tag{2.5}$$

which is the same as in a second-price auction (see (2.1)). Figure 2.3 depicts both the expected payment and the expected payoff of a bidder with value x in either auction. Because the expected revenue of the seller is just the sum of the *ex ante* (prior to knowing their values) expected payments of the bidders, this also implies that the expected revenues in the two auctions are the same. Let us see why.

The *ex ante* expected payment of a particular bidder in either auction is

$$\begin{aligned} E\left[m^A(X)\right] &= \int_0^{\omega} m^A(x) f(x)\, dx \\ &= \int_0^{\omega} \left(\int_0^{x} y g(y)\, dy \right) f(x)\, dx \end{aligned}$$

where $A = \mathrm{I}$ or II. Interchanging the order of integration, we obtain that

$$\begin{aligned} E\left[m^A(X)\right] &= \int_0^{\omega} \left(\int_y^{\omega} f(x)\, dx \right) y g(y)\, dy \\ &= \int_0^{\omega} y(1 - F(y))\, g(y)\, dy \end{aligned} \tag{2.6}$$

The expected revenue accruing to the seller $E\left[R^A\right]$ is just N times the *ex ante* expected payment of an individual bidder, so

$$\begin{aligned} E\left[R^A\right] &= N \times E\left[m^A(X)\right] \\ &= N \int_0^{\omega} y(1 - F(y))\, g(y)\, dy \end{aligned}$$

But now notice that the density of $Y_2^{(N)}$, the second highest of N values, $f_2^{(N)}(y) = N(1-F(y))f_1^{(N-1)}(y)$ (see Appendix C), and since $f_1^{(N-1)}(y) = g(y)$, we can write

$$\begin{aligned} E\left[R^A\right] &= \int_0^{\omega} y f_2^{(N)}(y)\, dy \\ &= E\left[Y_2^{(N)}\right] \end{aligned} \tag{2.7}$$

In either case, the expected revenue is just the expectation of the second-highest value. Thus, we conclude that *the expected revenues of the seller in the two auctions are the same*. For future reference, we record this fact in the following proposition.

Proposition 2.3. *With independently and identically distributed private values, the expected revenue in a first-price auction is the same as the expected revenue in a second-price auction.*

The fact that the expected selling prices in the two auctions are equal is all the more striking because in specific realizations of the values the price at which the object is sold may be greater in one auction or the other. With positive probability, the revenue R^{I} in a first-price auction exceeds R^{II}, the revenue in a second-price auction, and vice versa. For instance, when values are uniformly distributed and there are only two bidders, the equilibrium strategy in a first-price auction is $\beta^{\mathrm{I}}(x) = \frac{1}{2}x$. If the realized values are such that $\frac{1}{2}x_1 > x_2$, then the revenue in a first-price auction is greater than that in a second-price auction. On the other hand, if $\frac{1}{2}x_1 < x_2 < x_1$, the opposite is true. Thus, while the revenue may be greater in one auction or another depending on the realized values, we have argued that *on average* the revenue to the seller will be the same.

Actually, we can say more about the distribution of prices in the two auctions. It is clear that the revenues in a second-price auction are more variable than in its first-price counterpart. In the former, the prices can range between 0 and ω; in the latter, they can only range between 0 and $E[Y_1]$. A more precise result can be formulated along the following lines. Let L^{I} denote the distribution of the equilibrium price in a first-price auction and likewise, let L^{II} be the distribution of prices in a second-price auction. Then L^{II} is a *mean-preserving spread* of L^{I}—from the perspective of the seller, a second-price auction is *riskier* than a first-price auction (see Appendix B). Every risk-averse seller prefers the latter to the former (assuming, of course, that bidders are risk-neutral).[1] Figure 2.4 depicts the two distributions in the case of uniformly distributed values with two bidders. Since the two distributions have the same mean, the two shaded regions are, as they must be, equal in area.

---

[1]This is also equivalent to the statement that L^{I} dominates L^{II} in the sense of *second*-order stochastic dominance. Again, see Appendix B.

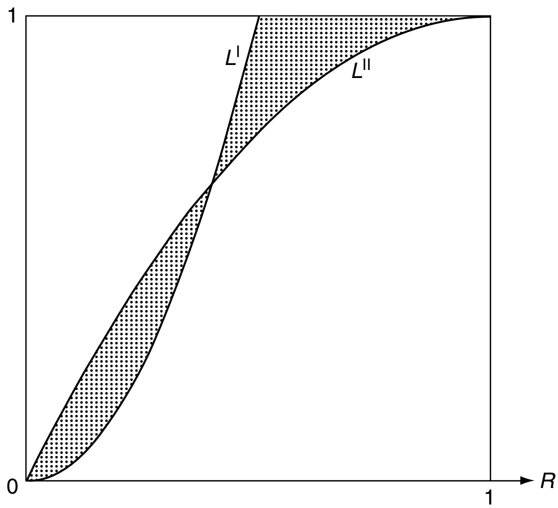

FIGURE 2.4 Distribution of prices in first- and second-price auctions.

Proposition 2.4. *With independently and identically distributed private values, the distribution of equilibrium prices in a second-price auction is a mean-preserving spread of the distribution of equilibrium prices in a first-price auction.*

Proof. The revenue in a second-price auction is just the random variable $R^{\mathrm{II}} = Y_2^{(N)}$; the revenue in a first-price auction is the random variable $R^{\mathrm{I}} = \beta\big(Y_1^{(N)}\big)$, where $\beta \equiv \beta^{\mathrm{I}}$ is the symmetric equilibrium strategy from Proposition 2.2. So we can write

$$E\left[R^{\mathrm{II}} \mid R^{\mathrm{I}} = p\right] = E\left[Y_2^{(N)} \mid Y_1^{(N)} = \beta^{-1}(p)\right]$$

But for all y,

$$E\left[Y_2^{(N)} \mid Y_1^{(N)} = y\right] = E\left[Y_1^{(N-1)} \mid Y_1^{(N-1)} < y\right] \tag{2.8}$$

This is because the only information regarding the second highest of N values, $Y_2^{(N)}$, that the event that the highest of N values $Y_1^{(N)} = y$ provides is that the highest of $N-1$ values, $Y_1^{(N-1)}$, is less than y. (See (C.6) in Appendix C for a formal demonstration.)

Using (2.8), we can write

$$\begin{aligned} E\left[R^{\mathrm{II}} \mid R^{\mathrm{I}} = p\right] &= E\left[Y_1^{(N-1)} \mid Y_1^{(N-1)} < \beta^{-1}(p)\right] \\ &= \beta\left(\beta^{-1}(p)\right) \\ &= p \end{aligned}$$

recalling (2.4).

Since $E\left[R^{\mathrm{II}} \mid R^{\mathrm{I}}=p\right]=p$, there exists a random variable Z such that the distribution of R^{II} is the same as that of $R^{\mathrm{I}}+Z$ and $E\left[Z \mid R^{\mathrm{I}}=p\right]=0$. Thus, L^{II} is a mean-preserving spread of L^{I}. ■

2.5 RESERVE PRICES

In the analysis so far, the seller has played a passive role. Indeed, we have implicitly assumed that the seller parts with the object at whatever price it will fetch. In many instances, sellers reserve the right to not sell the object if the price determined in the auction is lower than some threshold amount—say, $r>0$. Such a price is called the *reserve price*. We now examine what effect such a reserve price has on the expected revenue accruing to the seller.

RESERVE PRICES IN SECOND-PRICE AUCTIONS

Suppose that the seller sets a "small" reserve price of $r>0$. Since the price at which the object is sold can never be lower than r, no bidder with a value $x<r$ can make a positive profit in the auction. In a second-price auction, a reserve price makes no difference to the behavior of the bidders; it is still a weakly dominant strategy to bid one's value. The expected payment of a bidder with value r is now just $rG(r)$, and the expected payment of a bidder with value $x \geq r$ is

$$m^{\mathrm{II}}(x,r)=rG(r)+\int_r^x yg(y)\,dy \tag{2.9}$$

since the winner pays the reserve price r whenever the second-highest bid is below r.

RESERVE PRICES IN FIRST-PRICE AUCTIONS

Now consider a first-price auction with a reserve price $r>0$. Once again, since the price is at least r, no bidder with a value $x<r$ can make a positive profit. Furthermore, if β^{I} is a symmetric equilibrium of the first-price auction with reserve price r, it must be that $\beta^{\mathrm{I}}(r)=r$. This is because a bidder with value r wins only if all other bidders have values less than r and, in that case, can win with a bid of r itself. In all other respects, the analysis of a first-price auction is unaffected, and in a manner analogous to Proposition 2.2 we obtain that a symmetric equilibrium bidding strategy for any bidder with value $x \geq r$ is

$$\begin{aligned}
\beta^{\mathrm{I}}(x) &= E\left[\max\{Y_1,r\} \mid Y_1<x\right] \\
&= r\frac{G(r)}{G(x)}+\frac{1}{G(x)}\int_r^x yg(y)\,dy
\end{aligned}$$

The expected payment of a bidder with value $x \geq r$ is

$$\begin{aligned} m^{\mathrm{I}}(x,r) &= G(x) \times \beta^{\mathrm{I}}(x) \\ &= rG(r) + \int_r^x yg(y)\,dy \end{aligned} \tag{2.10}$$

which is the same as in (2.9).

Thus, once again, the expected payments and hence the expected revenues in the first- and second-price auctions are the same. Proposition 2.3 generalizes so as to accommodate reserve prices.

REVENUE EFFECTS OF RESERVE PRICES

How do reserve prices affect the seller's expected revenue? As before, let A denote either the first- or second-price auction. In both, the expected payment of a bidder with value r is $rG(r)$. A calculation similar to that in (2.6) shows that the *ex ante* expected payment of a bidder is now

$$\begin{aligned} E\left[m^A(X,r)\right] &= \int_r^{\omega} m^A(x,r) f(x)\,dx \\ &= r(1-F(r))\,G(r) + \int_r^{\omega} y(1-F(y))\,g(y)\,dy \end{aligned}$$

What is the optimal, or revenue maximizing, reserve price from the perspective of the seller? Suppose that the seller attaches a value $x_0 \in [0,\omega)$. This means that if the object is left unsold, the seller would derive a value x_0 from its use. Clearly, the seller would not set a reserve price r that is lower than x_0. Then the overall expected payoff of the seller from setting a reserve price $r \geq x_0$ is

$$\Pi_0 = N \times E\left[m^A(X,r)\right] + F(r)^N x_0$$

Differentiating this with respect to r, we obtain

$$\frac{d\Pi_0}{dr} = N[1-F(r) - rf(r)]\,G(r) + NG(r)f(r)x_0$$

Now recall that the *hazard rate* function associated with the distribution F is defined as $\lambda(x) = f(x)/(1-F(x))$. Thus, we can write

$$\frac{d\Pi_0}{dr} = N[1-(r-x_0)\lambda(r)]\,(1-F(r))\,G(r) \tag{2.11}$$

First, notice that if $x_0 > 0$, then the derivative of Π_0 at $r = x_0$ is positive, implying that the seller should set a reserve price $r > x_0$. If $x_0 = 0$, then the derivative of Π_0 at $r = 0$ is 0, but as long as $\lambda(r)$ is bounded, the expected payment attains a local minimum at 0, so a small reserve price leads to an increase in revenue. Thus, *a revenue maximizing seller should always set a reserve price that exceeds his or her value*. Why does a reserve price that exceeds x_0 lead to an increase in revenue? Consider a second-price auction with two bidders and suppose $x_0 = 0$. By setting a positive reserve price r the seller runs the risk that if the highest value among the bidders, Y_1, is smaller than r, the object will remain unsold. This potential loss is offset, however, by the possibility that while the highest-value Y_1 exceeds r, the second-highest value, Y_2, is smaller than r (in all other cases, the reserve price has no effect). Now the application of the reserve price means that the object will be sold for r rather than Y_2. The probability of the first event is $F(r)^2$ and the loss is at most r. So for small r, the expected loss is at most $rF(r)^2$. The probability of the second event is $2F(r)(1-F(r))$, and for small r, the gain is of order r, so the expected gain is of order $2rF(r)(1-F(r))$. Thus, the expected gain from setting a small reserve price exceeds the expected loss. This fact is sometimes referred to as the *exclusion principle*, since it implies, in effect, that it is optimal for the seller to exclude some bidders—those with values below the reserve price—from the auction even though their values exceed x_0.

Second, the relevant first-order condition implies that the optimal reserve price r^* must satisfy

$$\left(r^* - x_0\right)\lambda\left(r^*\right) = 1$$

or equivalently,

$$r^* - \frac{1}{\lambda\left(r^*\right)} = x_0 \tag{2.12}$$

If $\lambda(\cdot)$ is increasing, this condition is also sufficient, and it is remarkable that the optimal reserve price does not depend on the number of bidders. Roughly, the reason is that a reserve price comes into play only in instances when there is a single bidder with a value that exceeds the reserve price. So when a marginal change in the reserve price matters, it affects revenues in the same way as if there were a single bidder. Figure 2.5 depicts the expected revenue as a function of the reserve price r when F is the uniform distribution on $[0, 1]$, there are only two bidders, and $x_0 = 0$. As is clear from the figure, the optimal reserve price $r^* = \frac{1}{2}$. The resulting expected revenue is $\frac{5}{12}$.

ENTRY FEES

A positive reserve price r results in bidders with low values, lying below r, being excluded from the auction. Since their equilibrium payoffs are zero, such bidders are indifferent between participating in the auction and not. An alternative instrument that the seller can also use to exclude buyers with low values is an *entry fee*—a fixed and nonrefundable amount that bidders must pay the seller prior to the auction in order to be able to submit bids. An entry fee is, as it were, the price of admission to the room in which the auction is being conducted.

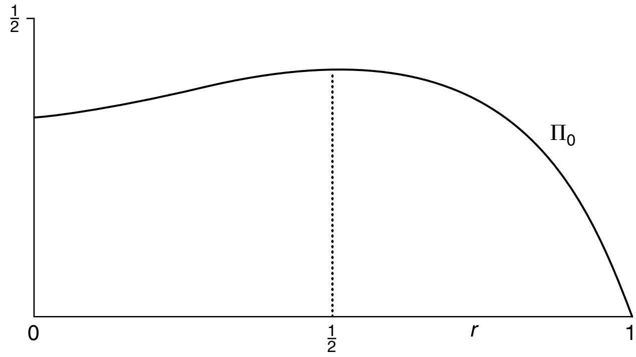

FIGURE 2.5 Optimal reserve price with uniformly distributed values.

A reserve price of r excludes all bidders with values $x < r$. The same set of bidders can be excluded by asking each bidder to pay an entry fee $e = G(r) \times r$. Notice that after paying e, the expected payoff of a bidder with value r in either a first- or second-price auction is exactly zero, so a bidder with value $x < r$ would not find it worthwhile to pay e in order to participate in the auction. The exclusion effect of a reserve price r can be replicated with an entry fee of e as determined earlier. Conversely, the exclusion effect of an entry fee e can be duplicated with a reserve price of r, again, as determined earlier.

EFFICIENCY VERSUS REVENUE

A reserve price (or equivalently, an entry fee) raises the revenue to the seller but may have a detrimental effect on efficiency. Suppose that the value that the seller attaches to the object is 0. In the absence of a reserve price, the object will always be sold to the highest bidder and in the symmetric model studied here, that is also the bidder with the highest value. Thus, both the first- and second-price auctions allocate efficiently in the sense that the object ends up in the hands of the person who values it the most. If the seller sets a reserve price $r > 0$, however, there is a positive probability that the object will remain in the hands of the seller and this is inefficient. This simple observation implies that there may be a trade-off between efficiency and revenue.

COMMITMENT ISSUES

There are two practical considerations that we have neglected. First, we have implicitly assumed that the seller can credibly commit to not sell the object if it cannot be sold at or above the reserve price. This commitment is particularly important because by setting a reserve price the seller is giving up some gains from trade. Without such a commitment, buyers may anticipate that the object, if durable, will be offered for sale again in a later auction and perhaps with a lower reserve price. These expectations may affect their bidding behavior in the first auction. Indeed, in the absence of a credible "no sale" commitment, the problem confronting a seller is analogous to that of a durable goods monopoly.

In both, a potential future sale may cause buyers to wait for lower prices, and this may reduce demand today. In effect, potential future sales may compete with current sales. In response, the seller may have to set lower reserve prices today than would be optimal in a one-time sale or if the good were perishable.

A second and not unrelated issue concerns secret reserve prices. We have assumed that the reserve price is publicly announced prior to the auction. In many situations, especially in art auctions, it is announced that there is a reserve price, but the level of the reserve price is not disclosed. In effect, the seller can opt to not sell the object after learning all the bids and thus the price. But this is rational only if the seller anticipates that in a future sale the price will be higher. Once again, buyers' expectations regarding future sales may affect the bidding in the current auction.

PROBLEMS

2.1. (Power distribution) Suppose there are two bidders with private values that are distributed independently according to the distribution $F(x) = x^a$ over $[0, 1]$, where $a > 0$. Find symmetric equilibrium bidding strategies in a first-price auction.

2.2. (Pareto distribution) Suppose there are two bidders with private values that are distributed independently according to a Pareto distribution $F(x) = 1 - (x+1)^{-2}$ over $[0, \infty)$. Find symmetric equilibrium bidding strategies in a first-price auction. Show by direct computation that the expected revenues in a first- and second-price auction are the same.

2.3. (Stochastic dominance) Consider an N-bidder first-price auction with independent private values. Let β be the symmetric equilibrium bidding strategy when each bidder's value is distributed according to F on $[0, \omega]$. Similarly, let β^* be the equilibrium strategy when each bidder's value distribution is F^* on $[0, \omega^*]$.

- **a.** Show that if F^* dominates F in terms of the reverse hazard rate (see Appendix B for a definition), then for all $x \in [0, \omega]$, $\beta^*(x) \geq \beta(x)$.
- **b.** By considering $F(x) = 3x - x^2$ on $\left[0, \frac{1}{2}(3 - \sqrt{5})\right]$ and $F^*(x) = 3x - 2x^2$ on $\left[0, \frac{1}{2}\right]$, show that the condition that F^* first-order stochastically dominates F is not sufficient to guarantee that $\beta^*(x) \geq \beta(x)$.

2.4. (Mixed auction) Consider an N-bidder auction which is a "mixture" of a first- and second-price auction in the sense that the highest bidder wins and pays a convex combination of his own bid and the second-highest bid. Precisely, there is a fixed $\alpha \in (0, 1)$ such that upon winning, bidder i pays $\alpha b_i + (1 - \alpha)\left(\max_{j \neq i} b_j\right)$. Find a symmetric equilibrium bidding strategy in such an auction when all bidders' values are distributed according to F.

2.5. (Resale) Consider a two-bidder first-price auction in which bidders' values are distributed according to F. Let β be the symmetric equilibrium (as derived in Proposition 2.2). Now suppose that after the auction is over, both the losing and winning bids are publicly announced. In addition, there is the possibility of postauction *resale*: The winner of the auction may, if he so wishes, offer the object to the other bidder at a fixed "take-it-or-leave-it" price of p. If the other bidder agrees, then the object changes hands, and the losing bidder pays the winning bidder p. Otherwise, the object stays with the winning bidder, and no money changes hands. The possibility of postauction resale in this manner is commonly known to both bidders prior to participating in the auction. Show that β remains an equilibrium even if resale is allowed. In particular, show that a bidder with value x cannot gain by bidding an amount $b > \beta(x)$ even when he has the option of reselling the object to the other bidder.

CHAPTER NOTES

The basic model of auctions with independent private values was introduced by Vickrey (1961). He derived equilibrium bidding strategies in a first-price auction when values are drawn from the uniform distribution (Example 2.1) and observed that the expected revenues in the first- and second-price auctions were the same. He recognized that this equivalence held more generally—that is, for arbitrary distributions—and formally established this in a subsequent paper, Vickrey (1962).

The symmetric independent private values model was analyzed in more detail by Riley and Samuelson (1981). The treatment of reserve prices follows that in Myerson (1981) and in Riley and Samuelson (1981). The problem of commitment in connection with reserve prices is discussed in Milgrom (1987) and formally analyzed in McAfee and Vincent (1997).

Chapter | Three

The Revenue Equivalence Principle

In the previous chapter we saw that regardless of the distribution of values, the expected selling price in a symmetric first-price auction is the same as that in a second-price auction. As a result a risk-neutral seller is indifferent between the two formats. The fact that the expected selling prices in the two auctions are equal is quite remarkable. The two auctions are not strategically equivalent as defined in Chapter 1, and in particular instances, the price in one or the other auction may be higher. This chapter explores the reasons underlying the equality of expected revenues in Proposition 2.3. In the process, we will discover that this equality extends beyond first- and second-price auctions to a whole class of auction forms.

3.1 MAIN RESULT

The auction forms we consider all have the feature that buyers are asked to submit *bids*—amounts of money they are willing to pay. These bids alone determine who wins the object and how much the winner pays. We will say that an auction is *standard* if the rules of the auction dictate that the person who bids the highest amount is awarded the object. Both first- and second-price auctions are, of course, standard in this sense, but so, for instance, is a *third-price auction*, discussed later in this chapter, in which the winner is the person bidding the highest amount but pays the third-highest bid. An example of a nonstandard method is a *lottery* in which the chances that a particular bidder wins is the ratio of his or her bid to the total amount bid by all. Such a lottery is nonstandard, since the person who bids the most is not necessarily the one who is awarded the object.

Given a standard auction form, A, and a symmetric equilibrium β^A of the auction, let $m^A(x)$ be the equilibrium *expected payment* by a bidder with value x. It turns out, quite remarkably, that provided that the expected payment of a bidder with value 0 is 0, the expected payment function $m^A(\cdot)$ does not depend on the particular auction form A. As a result, the expected revenue in any standard auction is the same, a fact known as the revenue equivalence principle.

Proposition 3.1. *Suppose that values are independently and identically distributed and all bidders are risk neutral. Then any symmetric and increasing equilibrium of any standard auction, such that the expected payment of a bidder with value zero is zero, yields the same expected revenue to the seller.*

Proof. Consider a standard auction form, A, and fix a symmetric equilibrium β of A. Let $m^A(x)$ be the equilibrium expected payment in auction A by a bidder with value x. Suppose that β is such that $m^A(0) = 0$.

Consider a particular bidder—say, 1—and suppose other bidders are following the equilibrium strategy β. It is useful to abstract away from the details of the auction and consider the expected payoff of bidder 1 with value x and when he bids $\beta(z)$ instead of the equilibrium bid $\beta(x)$. Bidder 1 wins when his bid $\beta(z)$ exceeds the highest competing bid $\beta(Y_1)$, or equivalently, when $z > Y_1$. His expected payoff is

$$\Pi^A(z,x) = G(z)x - m^A(z)$$

where as before $G(z) \equiv F(z)^{N-1}$ is the distribution of Y_1. The key point is that $m^A(z)$ depends on the other players' strategy β and z but is independent of the true value, x.

Maximization results in the first-order condition,

$$\frac{\partial}{\partial z}\Pi^A(z,x) = g(z)x - \frac{d}{dz}m^A(z) = 0$$

At an equilibrium it is optimal to report $z = x$, so we obtain that for all y,

$$\frac{d}{dy}m^A(y) = g(y)y \tag{3.1}$$

Thus,

$$\begin{aligned} m^A(x) &= m^A(0) + \int_0^x y g(y)\, dy \\ &= \int_0^x y g(y)\, dy \\ &= G(x) \times E[Y_1 \mid Y_1 < x] \end{aligned} \tag{3.2}$$

since, by assumption, $m^A(0) = 0$. Since the right-hand side does not depend on the particular auction form A, this completes the proof. ■

For the specification in Example 2.1, the expected payment function can be easily calculated.

Example 3.1. *Values are uniformly distributed on* $[0,1]$.

If $F(x)=x$, then $G(x)=x^{N-1}$ and for any standard auction satisfying $m^A(0)=0$, (3.2) implies that

$$m^A(x)=\frac{N-1}{N}x^N$$

and

$$E\left[m^A(X)\right]=\frac{N-1}{N(N+1)}$$

while the expected revenue is

$$E\left[R^A\right]=N\times E\left[m^A(X)\right]=\frac{N-1}{N+1}$$

▲

3.2 SOME APPLICATIONS OF THE REVENUE EQUIVALENCE PRINCIPLE

The revenue equivalence principle is a powerful and useful tool. In this section we show how, with judicious use, it can be used to derive equilibrium bidding strategies in alternative, unusual auction forms. We then show how it can be extended and applied to situations in which bidders are unsure as to how many other, rival bidders they face.

3.2.1 Unusual Auctions

We consider two unusual formats: an all-pay auction and a third-price auction. Although neither is used as a real-world auction to sell objects, the former is a useful model of other auction-like contests (some examples are offered next), while the latter is a useful theoretical construct.

EQUILIBRIUM OF ALL-PAY AUCTIONS

Consider an *all-pay* auction with the following rules. Each bidder submits a bid, and, as in the standard auctions discussed earlier, the highest bidder wins the object. The unusual aspect of an all-pay auction is that all bidders pay what they bid. The all-pay auction is a useful model of lobbying activity. In such models, different interest groups spend money—their "bids"—in order to influence government policy and the group spending the most—the highest "bidder"—is able to tilt policy in its favored direction, thereby "winning the auction." Since money spent on lobbying is a sunk cost borne by all groups regardless of which group is successful in obtaining its preferred policy, such situations have a natural all-pay aspect. We are interested in symmetric equilibrium strategies in an all-pay auction with symmetric, independent private values.

Suppose for the moment that there is a symmetric, increasing equilibrium of the all-pay auction such that the expected payment of a bidder with value 0 is 0. In other words, the assumptions of Proposition 3.1 are satisfied. Then we know that the expected payment in such an equilibrium must be the same as in (3.2). Now in an all-pay auction, the expected payment of a bidder with value x is the *same* as his bid—he forfeits his bid regardless of whether he wins or not—and so if there is a symmetric, increasing equilibrium of the all-pay auction β^{AP}, it must be that

$$\begin{aligned} \beta^{\mathrm{AP}}(x) &= m^A(x) \\ &= \int_0^x y g(y)\, dy \end{aligned}$$

To verify that this indeed constitutes an equilibrium of the all-pay auction, suppose that all bidders except one are following the strategy $\beta \equiv \beta^{\mathrm{AP}}$. If he bids an amount $\beta(z)$, the expected payoff of a bidder with value x is

$$G(z)x - \beta(z) = G(z)x - \int_0^z y g(y)\, dy$$

and integrating the second term by parts, this becomes

$$G(z)\,(x - z) + \int_0^z G(y)\, dy$$

which is the same as the payoff obtained in a first-price auction by bidding $\beta^{\mathrm{I}}(z)$ against other bidders who are following β^{I}. For the same reasons as in Proposition 2.2, this is maximized by choosing $z = x$. Thus, β^{AP} is a symmetric equilibrium.

EQUILIBRIUM OF THIRD-PRICE AUCTIONS

Suppose that there are at least three bidders. Consider a sealed-bid auction in which the highest bidder wins the object but pays a price equal to the third-highest bid. A *third*-price auction, as it is called, is a purely theoretical construct: There is no known instance of such a mechanism actually being used. It is an interesting construct nevertheless; equilibria of such an auction display some unusual properties, and it leads to a better understanding of the workings of the standard auction forms. Here we show how the revenue equivalence principle can once again be used to derive equilibrium bidding strategies.

Again, suppose for the moment that there is a symmetric, increasing equilibrium of the third-price auction—say, β^{III}—such that the expected payment of a bidder with value 0 is 0. Once again, since the assumptions of Proposition 3.1 are satisfied, we must have that for all x, the expected payment in a third-price auction is

$$m^{\mathrm{III}}(x) = \int_0^x y g(y)\, dy \tag{3.3}$$

On the other hand, consider bidder 1, and suppose that he wins in equilibrium when his value is x. Winning implies, of course, that his value x exceeds the highest of the other $N-1$ values—that is, $Y_1 < x$. The price bidder 1 pays is the random variable $\beta^{\mathrm{III}}(Y_2)$, where Y_2 is the second highest of the $N-1$ other values. The density of Y_2, conditional on the event that $Y_1 < x$, can be written as

$$f_2^{(N-1)}(y \mid Y_1 < x) = \frac{1}{F_1^{(N-1)}(x)}(N-1)(F(x)-F(y))f_1^{(N-2)}(y),$$

where $(N-1)(F(x)-F(y))$ is the probability that Y_1 exceeds $Y_2 = y$ but is less than x and $f_1^{(N-2)}(y)$ is the density of the highest of $N-2$ values. The expected payment in a third-price auction can then be written as

$$\begin{aligned} m^{\mathrm{III}}(x) &= F_1^{(N-1)}(x)E\left[\beta^{\mathrm{III}}(Y_2) \mid Y_1 < x\right] \\ &= \int_0^x \beta^{\mathrm{III}}(y)(N-1)(F(x)-F(y))f_1^{(N-2)}(y)dy \end{aligned} \tag{3.4}$$

Equating (3.3) and (3.4), we obtain that

$$\int_0^x \beta^{\mathrm{III}}(y)(N-1)(F(x)-F(y))f_1^{(N-2)}(y)dy = \int_0^x yg(y)\,dy$$

and differentiating with respect to x, this implies that

$$\begin{aligned} (N-1)f(x)\int_0^x \beta^{\mathrm{III}}(y)f_1^{(N-2)}(y)dy &= xg(x) \\ &= x \times (N-1)f(x)F(x)^{N-2} \end{aligned}$$

since $G(x) \equiv F(x)^{N-1}$. This can be rewritten as

$$\int_0^x \beta^{\mathrm{III}}(y)f_1^{(N-2)}(y)dy = xF_1^{(N-2)}(x)$$

since, $F_1^{(N-2)}(x) \equiv F(x)^{N-2}$. Differentiating once more with respect to x,

$$\beta^{\mathrm{III}}(x)f_1^{(N-2)}(x) = xf_1^{(N-2)}(x) + F_1^{(N-2)}(x)$$

and rearranging this we get

$$\begin{aligned} \beta^{\mathrm{III}}(x) &= x + \frac{F_1^{(N-2)}(x)}{f_1^{(N-2)}(x)} \\ &= x + \frac{F(x)}{(N-2)f(x)} \end{aligned}$$

This derivation, however, is valid only if β^{III} is increasing, and from the preceding equation it is clear that a sufficient condition for this is that the ratio F/f is increasing. This condition is the same as requiring that $\ln F$ is a concave function or equivalently that F is *log-concave*.

Proposition 3.2. *Suppose that there are at least three bidders and F is log-concave. Symmetric equilibrium strategies in a third-price auction are given by*

$$\beta^{\mathrm{III}}(x) = x + \frac{F(x)}{(N-2)f(x)} \tag{3.5}$$

An important feature of the equilibrium in a third-price auction is worth noting: The equilibrium bid *exceeds* the value. To better understand this phenomenon, first notice that for much the same reason as in a second-price auction, it is dominated for a bidder to bid below his value in a third-price auction. Unlike in a second-price auction, however, it is not dominated for a bidder to bid above his value. Fix some equilibrium bidding strategies of the third-price auction—say, β—and suppose that all bidders except 1 follow β. Suppose bidder 1 with value x bids an amount $b > x$. If $\beta(Y_2) < x < \beta(Y_1) < b$, this is better than bidding x, since it results in a profit, whereas bidding x would not. If, however, $x < \beta(Y_2) < \beta(Y_1) < b$, then bidding b results in a loss. When $b - x \equiv \varepsilon$ is small, the gain in the first case is of order ε^2, whereas the loss in the second case is of order ε^3. Thus, it is optimal to bid higher than one's value in a third-price auction.

Comparing equilibrium bids in first-, second-, and third-price auctions in case of symmetric private values, we have seen that

$$\beta^{\mathrm{I}}(x) < \beta^{\mathrm{II}}(x) = x < \beta^{\mathrm{III}}(x)$$

(assuming, of course, that the distribution of values is log-concave).

3.2.2 Uncertain Number of Bidders

In our analysis so far, each bidder knows his or her own value but is uncertain about the values of others. All other aspects of the situation—the number of bidders, the distribution from which they draw their values—are assumed to be common knowledge. In many auctions—particularly in those of the sealed-bid variety—a bidder may be uncertain about how many other interested bidders there are. In this section we show how the standard model may be amended to include this additional uncertainty.

Let $\mathcal{N} = \{1, 2, \ldots, N\}$ denote the set of *potential* bidders and let $\mathcal{A} \subseteq \mathcal{N}$ be the set of *actual* bidders—that is, those that participate in the auction. All potential bidders draw their values independently from the same distribution F.

Consider an actual bidder $i \in \mathcal{A}$ and let p_n denote the probability that any participating bidder assigns to the event that he is facing n other bidders. Thus, bidder i assigns the probability p_n that the number of actual bidders is $n+1$. The exact process by which the set of actual bidders is determined from the set of potential bidders is not important. What is important is that the process be symmetric so every actual bidder holds the *same* beliefs about how many other bidders he faces; the probabilities p_n do not depend on the identity of the bidder nor on his value. It is also important that the set of actual bidders does not depend on the realized values.

As long as bidders hold the same beliefs about the likelihood of meeting different numbers of rivals, the conclusion of Proposition 3.1 obtains in a straightforward manner. Consider a standard auction A and a symmetric and increasing equilibrium β of the auction. Note that since bidders are unsure about how many rivals they face, β does not depend on n. Consider the expected payoff of a bidder with value x who bids $\beta(z)$ instead of the equilibrium bid $\beta(x)$. The probability that he faces n other bidders is p_n. In that case, he wins if $Y_1^{(n)}$, the highest of n values drawn from F, is less than z and the probability of this event is $G^{(n)}(z) = F(z)^n$. The overall probability that he will win when he bids $\beta(z)$ is therefore

$$G(z) = \sum_{n=0}^{N-1} p_n G^{(n)}(z)$$

His expected payoff from bidding $\beta(z)$ when his value is x is then

$$\Pi^A(z,x) = G(z)x - m^A(z)$$

and the remainder of the argument is the same as in Proposition 3.1. Thus, we conclude that the revenue equivalence principle holds even if there is uncertainty about the number of bidders.

Suppose that the object is sold using a second-price auction. Even though the number of rival buyers that a particular bidder faces is uncertain, it is still a dominant strategy for him to bid his value. The expected payment in a second-price auction of an actual bidder with value x is therefore

$$m^{\mathrm{II}}(x) = \sum_{n=0}^{N-1} p_n G^{(n)}(x)\, E\left[Y_1^{(n)} \mid Y_1^{(n)} < x\right]$$

Now suppose that the object is sold using a first-price auction and that β is a symmetric and increasing equilibrium. The expected payment of an actual bidder with value x is

$$m^{\mathrm{I}}(x) = G(x)\,\beta(x)$$

where $G(x)$ is as defined earlier. The revenue equivalence principle implies that for all x, $m^{\mathrm{I}}(x) = m^{\mathrm{II}}(x)$, so

$$\begin{aligned}\beta(x) &= \sum_{n=0}^{N-1} \frac{p_n G^{(n)}(x)}{G(x)} E\left[Y_1^{(n)} \mid Y_1^{(n)} < x\right] \\ &= \sum_{n=0}^{N-1} \frac{p_n G^{(n)}(x)}{G(x)} \beta^{(n)}(x)\end{aligned}$$

where $\beta^{(n)}$ is the equilibrium bidding strategy in a first-price auction in which there are exactly $n+1$ bidders for sure (see Proposition 2.2 on page 15). Thus, the equilibrium bid for an actual bidder with value x when he is unsure about the number of rivals he faces is a weighted average of the equilibrium bids in auctions when the number of bidders is known to all.

PROBLEMS

3.1. (War of attrition) Consider a two-bidder *war of attrition* in which the bidder with the highest bid wins the object but both bidders pay the losing bid. Bidders' values independently and identically distributed according to F.

- **a.** Use the revenue equivalence principle to derive a symmetric equilibrium bidding strategy in the war of attrition.
- **b.** Directly compute the symmetric equilibrium bidding strategy and the seller's revenue when the bidders' values are uniformly distributed on $[0, 1]$.

3.2. (Losers-pay auction) Consider a N-bidder *losers-pay auction* in which the bidder with the highest bid wins the object and pays nothing, while all losing bidders pays their own bids. Bidders' valuations independently and identically distributed according to F.

- **a.** Use the revenue equivalence principle to derive a symmetric equilibrium bidding strategy in the losers-pay auction.
- **b.** Directly compute the symmetric equilibrium bidding strategy for the case when the bidders' values are distributed according to $F(x) = 1 - e^{-ax}$ over $[0, \infty)$.

CHAPTER NOTES

The revenue equivalence principle was established by Riley and Samuelson (1981) and Myerson (1981), showing, in effect, that the phenomenon noticed by Vickrey (1961, 1962) was quite general.

For a model of interest group lobbying as an all-pay auction, albeit in a complete information setting, see Baye, Kovenock, and de Vries (1993). Third-price auctions were first analyzed by Kagel and Levin (1993) who pointed out the over-bidding phenomenon . The explicit derivation of equilibrium strategies is due to Wolfstetter (2001).

Auctions with an uncertain number of bidders have been considered by McAfee and McMillan (1987b), Matthews (1987), and Harstad, Kagel, and Levin (1990). The first two papers are particularly interested in how risk-averse bidders—considered in the next chapter—are affected by uncertainty regarding the number of competitors they face. Harstad, Kagel, and Levin (1990) derive equilibrium bidding strategies in different auctions under number uncertainty when bidders are risk neutral.

Chapter | Four

Qualifications and Extensions

The revenue equivalence principle derived in the previous chapter, Proposition 3.1, is a simple yet powerful result. It constitutes a benchmark of the theory of private value auctions, whereas all other results in the area constitute a departure from the revenue equivalence principle and can be measured against it. Because of its central nature, it is worthwhile to recount the key assumptions underlying the principle:

1. *Independence*—the values of different bidders are independently distributed.
2. *Risk neutrality*—all bidders seek to maximize their expected profits.
3. *No budget constraints*—all bidders have the ability to pay up to their respective values.
4. *Symmetry*—the values of all bidders are distributed according to same distribution function F.

In this chapter we investigate how the revenue equivalence principle is affected when some of these assumptions are relaxed. We first explore the consequences of risk aversion on the part of bidders. We then study the effects of the assumption that bidders have sufficient financial resources to pay any price up to their values. We ask how the revenue equivalence principle holds up in an augmented model in which bidders are subject to budget constraints. Finally, we delve into the important issue of *ex ante* heterogeneity among the bidders. In each case, to isolate the effects of each assumption, we retain the others. For instance, we examine the consequences of risk aversion, retaining the assumptions regarding the independence of values, the lack of budget constraints, and symmetry among the bidders. In the same vein, we examine the consequences of budget constraints in a model with risk-neutral, symmetric bidders with independently distributed values, and we explore the consequences of bidder asymmetries, retaining the independence of the values, and the risk neutrality of the bidders.

An exploration of the consequences of relaxing the first assumption—the independence of the values—is postponed for the moment. It is the focus of Chapter 6, where we consider a more general model that simultaneously relaxes both this assumption and the assumption of private values.

4.1 RISK-AVERSE BIDDERS

We now argue that if bidders are risk-averse, but all other assumptions are retained, the revenue equivalence principle no longer holds. In particular, we retain all our other assumptions: independence of values, symmetry among bidders, and the absence of budget constraints.

Risk neutrality implies that the expected payoff of a bidder is additively separable, it is just the difference between the bidder's expected gain and his expected payment, so the payoff is linear in the payments. This quasi-linearity of a bidder's payoff is crucial in the derivation of the revenue equivalence result and is lost when bidders are not risk neutral.

To examine the consequences of risk aversion, suppose that each bidder has a von-Neumann-Morgenstern utility function $u : \mathbb{R}_+ \to \mathbb{R}$ that satisfies $u(0) = 0$, $u' > 0$ and $u'' < 0$. Each bidder now seeks to maximize his or her expected utility rather than expected profits. The main finding is as follows:

Proposition 4.1. *Suppose that bidders are risk-averse with the same utility function. With symmetric, independent private values, the expected revenue in a first-price auction is greater than that in a second-price auction.*

Proof. First, notice that risk aversion makes no difference in a second-price auction: it is still a dominant strategy for each bidder to bid his or her value. Thus, in a second-price auction, the expected price is the same as it would be if bidders were risk neutral.

Let us now examine a first-price auction. Suppose that when bidders are risk averse and have the utility function u, the equilibrium strategies are given by an increasing and differentiable function $\gamma : [0, \omega] \to \mathbb{R}_+$ satisfying $\gamma(0) = 0$. If all but bidder 1, say, follow this strategy, then bidder 1 will never bid more than $\gamma(\omega)$. Given a value x, each bidder's problem is to choose $z \in [0, \omega]$ and bid an amount $\gamma(z)$ to maximize his or her expected utility—that is,

$$\max_z G(z) u(x - \gamma(z)) \tag{4.1}$$

where, as before, $G \equiv F^{N-1}$ is the distribution of the highest of $N - 1$ values. The first-order condition for this problem is

$$g(z) \times u(x - \gamma(z)) - G(z) \times \gamma'(z) \times u'(x - \gamma(z)) = 0$$

In a symmetric equilibrium, it must be optimal to choose $z = x$. Hence, we get

$$\frac{g(x) u(x - \gamma(x))}{\gamma'(x)} = G(x) u'(x - \gamma(x))$$

which is the same as

$$\gamma'(x) = \frac{u(x-\gamma(x))}{u'(x-\gamma(x))} \times \frac{g(x)}{G(x)} \tag{4.2}$$

With risk neutrality, $u(x) = x$, and (4.2) then yields

$$\beta'(x) = (x - \beta(x)) \times \frac{g(x)}{G(x)}, \tag{4.3}$$

where $\beta(\cdot)$ denotes the equilibrium strategy with risk-neutral bidders.

Next notice that if u is strictly concave and $u(0) = 0$, for all $y > 0$, $[u(y)/u'(y)] > y$. Using this fact, from (4.2) we can derive that

$$\gamma'(x) = \frac{u(x-\gamma(x))}{u'(x-\gamma(x))} \times \frac{g(x)}{G(x)} > (x - \gamma(x)) \times \frac{g(x)}{G(x)} \tag{4.4}$$

Now if $\beta(x) > \gamma(x)$, we have $(x - \gamma(x)) \times g(x)/G(x) > (x - \beta(x)) \times g(x)/G(x)$, and because of (4.4) this implies that $\gamma'(x) > \beta'(x)$.

To summarize, if $\beta(\cdot)$ and $\gamma(\cdot)$ are the equilibrium strategies with risk-neutral and risk-averse bidders, respectively,

$$\beta(x) > \gamma(x) \text{ implies that } \beta'(x) < \gamma'(x) \tag{4.5}$$

It is also easy to check that

$$\beta(0) = \gamma(0) = 0 \tag{4.6}$$

(4.5) and (4.6) imply that for all $x > 0$,

$$\gamma(x) > \beta(x)$$

Thus, in a first-price auction, risk aversion causes an increase in equilibrium bids. Since bids have increased, the expected revenue has also increased. Using Proposition 3.1 and the fact that the expected revenue in a second-price auction is unaffected by risk aversion, we deduce that the expected revenue in a first-price auction is higher than that in a second-price auction. ■

Why does risk aversion lead to higher bids in a first-price auction? Consider a particular bidder—say, 1—with value x. Fix the strategies of all other bidders, and then suppose bidder 1 bids the amount b. Now suppose that this bidder considers decreasing his bid slightly from b to $b - \Delta$. If he wins the auction with this lower bid, this leads to a gain of Δ. A lowering of his bid could, however, cause him to lose the auction. For a risk-averse bidder, the effect of a slightly lower winning bid on his wealth level has a smaller utility consequence than does the possible loss if this lower bid were, in fact, to result in his losing the auction. Compared to a risk-neutral bidder, a risk-averse bidder will thus bid higher. Put another way, by bidding higher, a risk-averse bidder will, as it were, "buy" insurance against the possibility of losing.

Example 4.1. *Constant relative risk aversion (CRRA) utility functions.*

Consider a situation with two bidders who display constant relative risk aversion: Their utility functions are of the form $u(z) = z^{\alpha}$, where α satisfies $0 < \alpha < 1$, so the coefficient of relative risk aversion, $-zu''(z)/u'(z)$, is $1 - \alpha$. Suppose that both values are drawn from the distribution F. It is convenient to define $F_{\alpha} \equiv F^{1/\alpha}$ and notice that F_{α} is also a distribution function with the same support as F. The symmetric equilibrium bidding strategy in a first-price auction is the solution to the differential equation in (4.2). With the specified utility function, (4.2) becomes

$$\gamma'(x)F(x) + \frac{1}{\alpha}\gamma(x)f(x) = \frac{1}{\alpha}xf(x)$$

together with the boundary condition $\gamma(0) = 0$. Using $F(x)^{(1/\alpha)-1}$ as the integrating factor, the solution to this is easily seen to be

$$\gamma(x) = \frac{1}{F_{\alpha}(x)}\int_0^x y f_{\alpha}(y)dy,$$

where $f_{\alpha} = F'_{\alpha}$. This, of course, is of the same form as derived in Proposition 2.2 on page 15.

Thus, we conclude that the equilibrium bidding strategy with two bidders with CRRA utility functions $u(z) = z^{\alpha}$ whose values are drawn from the distribution F is the *same* as the equilibrium bidding strategy with two risk-neutral bidders whose values are drawn from the distribution F_{α}. Since $F_{\alpha} \leq F$, the expected revenue in a first-price auction with risk-averse bidders is greater than with risk-neutral bidders. The expected revenue in a second-price auction is, of course, unchanged. ▲

Example 4.2. *Constant absolute risk aversion (CARA) utility functions.*

Consider a situation with bidders who exhibit constant absolute risk aversion: Their utility functions are of the form $u(z) = 1 - \exp(-\alpha z)$, where $\alpha > 0$ is the coefficient of absolute risk aversion, $-u''(z)/u'(z)$. Suppose that values are independently distributed according to the function F and let G denote, as usual, the distribution of the highest of $N - 1$ values. First, consider a second-price auction. Consider a bidder with value x who bids z and wins the auction. In a second-price auction, such a bidder faces some uncertainty about the price he or she will pay, since that is determined by the second-highest bid. Suppose that the other bidders are following their equilibrium (and weakly dominant) strategy of bidding their values so that the second-highest bid is Y_1. Let $\rho(x,z)$

be the *risk premium* associated with the "price gamble"; it is the certain amount the bidder would forgo in order to remove the associated uncertainty. Formally,

$$u(x-\rho(x,z)) = E\left[u(x-Y_1) \mid Y_1 < z\right] \tag{4.7}$$

and CARA implies that we can write $\rho(z) \equiv \rho(x,z)$, since the risk premium depends only on the gamble being faced—which is entirely determined by z—and not on the "wealth level" x. It is optimal for bidder 1 to bid his or her true value and thus

$$x \in \arg\max_z G(z) E\left[u(x-Y_1) \mid Y_1 < z\right]$$

Using (4.7) this can be rewritten as

$$x \in \arg\max_z G(z) u(x-\rho(z))$$

But this is the same as bidder 1's maximization problem in a *first*-price auction if all other bidders follow the bidding strategy $\gamma = \rho$ (see (4.1)). This implies that for CARA bidders, the equilibrium bidding strategy in a first-price auction is to bid the risk premium associated with "price gamble" in a second-price auction. Finally, since

$$G(x)u(x-\gamma(x)) = G(x)E\left[u(x-Y_1) \mid Y_1 < x\right]$$

the equilibrium expected utility of a CARA bidder in a first-price auction is the same as his expected utility in a second-price auction. ▲

A key feature of the standard auction model with risk-neutral bidders is that the payoff functions are separable in money. In particular, they are *quasi-linear*—linear in the payments that bidders make—and bidders maximize their expected profits, which are just

$$\text{Expected Value} - \text{Expected Payment}$$

This separation between the expected value and the expected payment is crucial for revenue equivalence principle. Specifically, in the proof of Proposition 3.1 on page 28, this separation leads to equation (3.1) and hence to the conclusion that the expected payments are the same in any standard auction. Risk-averse bidders, on the other hand, maximize

$$\text{Expected Utility of (Value} - \text{Payment)}$$

and since utility is concave, the maximand is no longer linear in the payments that bidders make. The fact that bidders' objective functions are no longer linear in their payments is the reason for the failure of the revenue equivalence principle.

4.2 BUDGET CONSTRAINTS

Until now we have implicitly assumed that bidders face no cash or credit constraints—that is, bidders are able to pay the seller up to amounts equal to their values. In many situations, however, bidders may face financial constraints of one sort or another. In this section we ask how the presence of financial constraints affects equilibrium behavior in first- and second-price auctions and what effect they have on the revenue from these auctions.

We continue with the basic symmetric independent private value setting introduced in the previous chapter. There is a single object for sale and N potential buyers are bidding for the object. As before, bidder i assigns a value of X_i to the object. But now, in addition, each bidder is subject to an absolute *budget* of W_i. In no circumstances can a bidder with value-budget pair (x_i, w_i) pay more than w_i. We also suppose that if bidder i were to bid more than w_i and *default*, then a (small) penalty would be imposed.

Each bidder's value-budget pair (X_i, W_i) is independently and identically distributed on $[0,1] \times [0,1]$ according to the density function f.[1] Bidder i knows the realized value-budget pair (x_i, w_i), and only that other bidders' budget-value pairs are independently distributed according to f. As before, bidders are assumed to be risk neutral, and again we compare first- and second-price auctions. In a substantive departure from the models studied so far, the private information of the bidders is two-dimensional. We will refer to the pair (x_i, w_i) as the *type* of bidder i.

In any auction A (say, a first- or second-price auction), a bidder's strategy is a function of the form $B^A : [0,1] \times [0,1] \to \mathbb{R}$ that determines the amount bid depending on both his value and his budget.

4.2.1 Second-Price Auctions

We begin our analysis by considering second-price auctions. In this case, bidders' equilibrium strategies are straightforward.

Proposition 4.2. *In a second-price auction, it is a dominant strategy to bid according to* $B^{\mathrm{II}}(x, w) = \min\{x, w\}$.

Proof. First, notice that it is dominated to bid above one's budget. Suppose bidder i wins by bidding above his budget. If the second-highest bid is below his budget, then he would have also won by bidding w_i. If the second-highest bid is above his budget, he has to renege, does not get the object, and pays the fine, resulting in a negative surplus.

Second, if $x_i \le w_i$, then the budget constraint does not bind and the same argument as in the unconstrained situation implies that it is a weakly dominant

---

[1]The independence holds only across bidders. The possibility that for each bidder the values and budgets are correlated is admitted.

strategy to bid x_i. If $x_i > w_i$, then the same argument shows that bidding w_i dominates bidding less. ■

For every type (x, w), define $x'' = \min\{x, w\}$ and consider the type $\left(x'', 1\right)$. Notice that since values never exceed 1, a bidder of type $\left(x'', 1\right)$ effectively never faces a financial constraint. But since $\min\{x'', 1\} = x'' = \min\{x, w\}$, we have that $B^{\mathrm{II}}(x, w) = B^{\mathrm{II}}(x'', 1)$. Thus, in a second-price auction the type $\left(x'', 1\right)$ would submit a bid identical to that submitted by type (x, w). The type $\left(x'', 1\right)$ is, as it were, the richest member of the family with types (x, w) such that $\min\{x, w\} = x''$. Figure 4.1 depicts the set of types who bid the same in a second-price auction as does type (x, w). This consists of all types on the thin-lined right angle "Leontief iso-bid" curve whose corner lies on the diagonal.

As before, let $m^{\mathrm{II}}(x, w)$ denote the expected payment of a bidder of type (x, w) in a second-price auction. Since $B^{\mathrm{II}}(x, w) = B^{\mathrm{II}}(x'', 1)$, we have that

$$m^{\mathrm{II}}(x, w) = m^{\mathrm{II}}(x'', 1) \tag{4.8}$$

Now define

$$\mathcal{L}^{\mathrm{II}}(x'') = \left\{(X, W) : B^{\mathrm{II}}(X, W) < B^{\mathrm{II}}(x'', 1)\right\} \tag{4.9}$$

to be the set of types who bid less than type $\left(x'', 1\right)$ in a second-price auction.

Define

$$F^{\mathrm{II}}(x'') = \int_{\mathcal{L}^{\mathrm{II}}(x'')} f(X, W)\, dX dW \tag{4.10}$$

to be the probability that a type $\left(x'', 1\right)$ will outbid *one* other bidder. Notice that this is indeed the distribution function of the random variable $X'' = \min\{X, W\}$. The probability that a type $\left(x'', 1\right)$ will actually win the auction is just $(F^{\mathrm{II}}(x''))^{N-1} \equiv G^{\mathrm{II}}(x'')$. In Figure 4.1, $F^{\mathrm{II}}(x'')$ is the probability mass attached to the set of types lying below the lighter of the two right angles.

Now notice that we can write the expected utility of a type $\left(x'', 1\right)$ when bidding $B^{\mathrm{II}}(z, 1)$ as

$$G^{\mathrm{II}}(z) x'' - m^{\mathrm{II}}(z, 1)$$

In equilibrium, it is optimal to bid $B^{\mathrm{II}}(x'', 1)$ when the true type is $\left(x'', 1\right)$, so in a manner completely analogous to the arguments in Chapter 3 (specifically, Proposition 3.1), we have that

$$m^{\mathrm{II}}(x'', 1) = \int_0^{x''} y g^{\mathrm{II}}(y) dy, \tag{4.11}$$

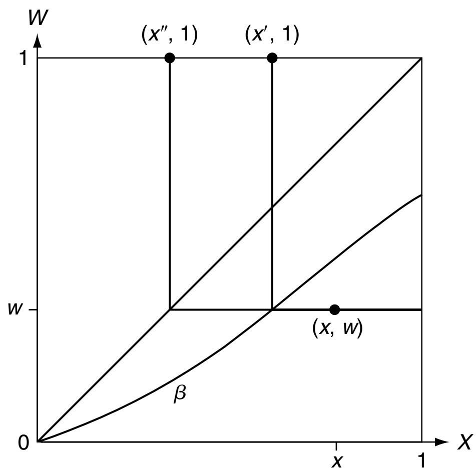

FIGURE 4.1 First- and second-price auctions with budget constraints.

where g^{II} is the density function associated with G^{II}. The *ex ante* expected payment of a bidder in a second-price auction with financial constraints can then be written, in manner completely analogous to (2.7), as

$$\begin{aligned} R^{\mathrm{II}} &= \int_0^1 m^{\mathrm{II}}(x'',1) f^{\mathrm{II}}(x'')\,dx'' \\ &= E\left[Y_2^{\mathrm{II}(N)}\right], \end{aligned} \tag{4.12}$$

where $Y_2^{\mathrm{II}(N)}$ is the second-highest of N draws from the distribution F^{II}.

4.2.2 First-Price Auctions

First, suppose that in a first-price auction, the equilibrium strategy is of the form

$$B^{\mathrm{I}}(x,w) = \min\{\beta(x), w\} \tag{4.13}$$

for some increasing function $\beta(x)$. In this case, it must be that $\beta(x) < x$, otherwise a bidder of type $x < w$ would deviate by bidding slightly less. Although sufficient conditions in terms of the primitives of the model that guarantee the existence of such an equilibrium can be provided, here we content ourselves with directly assuming that such an equilibrium exists.

As in the case of a second-price auction, for every type (x,w) define x' to be a value such that $\beta(x') = \min\{\beta(x), w\}$ and consider the type $(x',1)$. As before, a bidder of type $(x',1)$ effectively never faces a financial constraint. But since $\min\{\beta(x'),1\} = \beta(x') = \min\{\beta(x),w\}$, we have that $B^{\mathrm{I}}(x,w) = B^{\mathrm{I}}(x',1)$. Thus, in a first-price auction the type $(x',1)$ would submit a bid identical to that submitted by type (x,w). Figure 4.1 also depicts the set of types who bid the same in a first-price auction as does type (x, w). This consists of all types on the thick line right-angle "Leontief iso-bid" curve whose corner lies on the curve β.

Now define

$$\mathcal{L}^{\mathrm{I}}(x') = \left\{(X, W) : B^{\mathrm{I}}(X, W) < B^{\mathrm{I}}(x', 1)\right\} \tag{4.14}$$

and define m^{I}, F^{I}, and G^{I} in a fashion completely analogous to the corresponding objects for a second-price auction. Exactly the same reasoning shows that

$$E\left[R^{\mathrm{I}}\right] = E\left[Y_2^{\mathrm{I}(N)}\right] \tag{4.15}$$

where $Y_2^{\mathrm{I}(N)}$ is the second-highest of N draws from the distribution F^{I}.

4.2.3 Revenue Comparison

To compare the expected payments in the two auctions, notice that since for all x, $\beta(x) < x$, the definitions of $\mathcal{L}^{\mathrm{II}}(x)$ and $\mathcal{L}^{\mathrm{I}}(x)$ in (4.9) and (4.14), respectively, imply that $\mathcal{L}^{\mathrm{I}}(x) \subset \mathcal{L}^{\mathrm{II}}(x)$. See Figure 4.1. Now (4.10) implies that for all x, $F^{\mathrm{I}}(x) \leq F^{\mathrm{II}}(x)$ and a strict inequality holds for all $x \in (0, 1)$. We have thus argued that F^{I} stochastically dominates F^{II}. This implies that

$$E\left[Y_2^{\mathrm{I}(N)}\right] > E\left[Y_2^{\mathrm{II}(N)}\right]$$

Thus, we have shown the following:

Proposition 4.3. *Suppose that bidders are subject to financial constraints. If the first-price auction has a symmetric equilibrium of the form* $B^{\mathrm{I}}(x, w) = \min\{\beta(x), w\}$*, then the expected revenue in a first-price auction is greater than the expected revenue in a second-price auction.*

At an intuitive level, Proposition 4.3 results from the fact that budget constraints are "softer" in first-price auctions than in second-price auctions. Given a situation with budget constraints, consider a hypothetical situation in which each bidder's value is $Z_i = \min\{X_i, W_i\}$ and there are no budget constraints. Since value-budget pairs are independently and identically distributed across bidders, the revenue equivalence principle applies to the hypothetical situation, so the second-price and first-price auctions yield the same expected revenue—say, R. Now returning to the original situation with budget constraints, Proposition 4.2 implies that the revenue from the second-price auction in the original situation is also R; bids in the two are identical for every realization. The revenue in a first-price auction is greater than R because the comparison is between a situation in which bidders have values $X_i \geq Z_i$ and budgets $W_i \geq Z_i$ and a hypothetical situation in which they have values Z_i but no budget constraints.

4.3 ASYMMETRIES AMONG BIDDERS

In this section, we consider situations in which bidders are *ex ante* asymmetric: different bidders' values are drawn from different distributions. Asymmetries among bidders do not affect bidding behavior in second-price auctions; it is still a weakly dominant strategy for each bidder to bid his or her value. In a first-price auction, however, asymmetries lead to numerous complications. First, although an equilibrium exists (see Appendix G), unlike in the case of symmetric bidders, a closed form expression for the bidding strategies is not available, making a comparison with the second-price auction rather difficult. Second, the allocations under the two auctions are quite different: The second-price auction is efficient, whereas the first-price auction is not—and as a result, the two are no longer revenue equivalent. Indeed, as we will see, no general ranking of the revenues can be obtained.

We begin by exploring the nature of equilibrium bidding behavior in first-price auctions. To keep the analysis at a relatively simple level, we concentrate on the case of two bidders.

4.3.1 Asymmetric First-Price Auctions with Two Bidders

Suppose there are two bidders with values X_1 and X_2, which are independently distributed according to the functions F_1 on $[0,\omega_1]$ and F_2 on $[0,\omega_2]$, respectively. Suppose for the moment that there is an equilibrium of the first-price auction in which the two bidders follow the strategies β_1 and β_2, respectively. Further suppose that these are increasing and differentiable and have inverses $\phi_1 \equiv \beta_1^{-1}$ and $\phi_2 \equiv \beta_2^{-1}$, respectively.

It is clear that $\beta_1(0) = 0 = \beta_2(0)$, since it would be dominated for a bidder to bid more than the value. Moreover, $\beta_1(\omega_1) = \beta_2(\omega_2)$ since otherwise, if say, $\beta_1(\omega_1) > \beta_2(\omega_2)$, then bidder 1 would win with probability 1 when his value is ω_1 and would pay more than he needs to—he could increase his payoff by bidding slightly less than $\beta_1(\omega_1)$. Let

$$\overline{b} \equiv \beta_1(\omega_1) = \beta_2(\omega_2) \tag{4.16}$$

be the common highest bid submitted by either bidder.

Given that bidder $j = 1,2$ is following the strategy β_j, the expected payoff of bidder $i \neq j$ when his value is x_i and he bids an amount $b < \overline{b}$ is

$$\begin{aligned}\Pi_i(b,x_i) &= F_j\big(\phi_j(b)\big)\,(x_i - b)\\ &= H_j(b)\,(x_i - b)\end{aligned}$$

where $H_j(\cdot) \equiv F_j\big(\phi_j(\cdot)\big)$ denotes the distribution of bidder j's bids.

The first-order condition for bidder i requires that for all $b < \overline{b}$,

$$h_j(b)\,(\phi_i(b) - b) = H_j(b) \tag{4.17}$$

where $j \neq i$ and, as usual, $h_j(b) \equiv H_j'(b) = f_j(\phi_j(b))\,\phi_j'(b)$ is the density of j's bids. This can be rearranged so that

$$\phi_j'(b) = \frac{F_j(\phi_j(b))}{f_j(\phi_j(b))} \frac{1}{(\phi_i(b) - b)} \tag{4.18}$$

A solution to the system of differential equations in (4.18)—one for each bidder—together with the relevant boundary conditions constitutes an equilibrium of the first-price auction. Unfortunately, an explicit solution can be obtained only in some special cases—an example is given later—and so instead, we deduce some properties of the equilibrium strategies indirectly. To do this, we make some assumptions regarding the specific nature of the asymmetries.

WEAKNESS LEADS TO AGGRESSION

Suppose that bidder 1's values are "stochastically higher" than those of bidder 2. In particular, we will make the stronger assumption that the distribution F_1 *dominates* F_2 *in terms of the reverse hazard rate*—that is, $\omega_1 \geq \omega_2$ and for all $x \in (0, \omega_2)$,

$$\frac{f_1(x)}{F_1(x)} > \frac{f_2(x)}{F_2(x)} \tag{4.19}$$

Reverse hazard rate dominance is further discussed in Appendix B where it is also shown that it implies that F_1 stochastically dominates F_2—that is, $F_1(x) \leq F_2(x)$. If (4.19) holds, we will call bidder 1 the "strong" bidder and bidder 2 the "weak" bidder. (A simple class of examples in which the two distributions can be ordered according to the reverse hazard rate consists of distributions satisfying $F_1(x) = (F_2(x))^\theta$ for some $\theta > 1$.)

We now show that the weak bidder will bid more aggressively than the strong bidder in the sense that for any fixed value, the bid of the weak bidder will be higher than the bid of the strong bidder.

Proposition 4.4. *Suppose that the value distribution of bidder* 1 *dominates that of bidder* 2 *in terms of the reverse hazard rate. Then in a first-price auction, the "weak" bidder* 2 *bids more aggressively than the "strong" bidder* 1*—that is, for any* $x \in (0, \omega_2)$,

$$\beta_1(x) < \beta_2(x)$$

Proof. First, notice that if there exists a c such that $0 < c < \overline{b}$ and $\phi_1(c) = \phi_2(c) \equiv z$, then (4.18) and (4.19) imply that

$$\phi_2'(c) = \frac{F_2(z)}{f_2(z)} \frac{1}{(z-c)} > \frac{F_1(z)}{f_1(z)} \frac{1}{(z-c)} = \phi_1'(c)$$

Since $\phi_i'(c) = 1/\beta_i'(z)$, this is equivalent to saying that if there exists a z such that $\beta_1(z) = \beta_2(z)$, then $\beta_1'(z) > \beta_2'(z)$. In other words, if the curves β_1 and β_2 ever intersect, the former is steeper than the latter and this implies that they intersect at most once.

We will argue by contradiction. So suppose that there exists an $x \in (0, \omega_2)$ such that $\beta_1(x) \geq \beta_2(x)$. Then, either β_1 and β_2 do not intersect at all so $\beta_1 > \beta_2$ everywhere, or they intersect only once at some value $z \in (0, \omega_2)$ and for all x such that $z < x < \omega_2$, $\beta_1(x) > \beta_2(x)$. In either case, for all x close to ω_2, $\beta_1(x) > \beta_2(x)$.

Now notice that if $\omega_1 > \omega_2$, then from (4.16) $\beta_1(\omega_1) = \beta_2(\omega_2)$, so $\beta_1(\omega_2) < \beta_2(\omega_2)$. This contradicts the fact that for all x close to ω_2, $\beta_1(x) > \beta_2(x)$.

Next suppose $\omega_1 = \omega_2 \equiv \omega$. If we write $\beta_1(\omega) = \beta_2(\omega) = \overline{b}$, then in terms of the inverse bidding strategies we have that for all b close to $\overline{b}$, $\phi_1(b) < \phi_2(b)$. This implies that for all b close to $\overline{b}$,

$$H_1(b) = F_1(\phi_1(b)) \leq F_2(\phi_2(b)) = H_2(b)$$

and since $H_1\left(\overline{b}\right) = 1 = H_2\left(\overline{b}\right)$, it must be that $h_1(b) > h_2(b)$. Now using (4.17) we obtain that for all b close enough to $\overline{b}$,

$$\phi_1(b) = \frac{H_2(b)}{h_2(b)} + b > \frac{H_1(b)}{h_1(b)} + b = \phi_2(b),$$

which is a contradiction. ■

We know that F_1 stochastically dominates F_2 so that bidder 1's values are stochastically higher. At the same time, Proposition 4.4 shows that for any given value, bidder 2 bids higher than does bidder 1. What can be said about the distributions of bids, H_1 and H_2? Notice that since for all $b \in (0, \overline{b})$, $\phi_1(b) > \phi_2(b)$, it now follows from (4.17) and (4.18) that

$$\frac{H_2(b)}{h_2(b)} = \phi_1(b) - b > \phi_2(b) - b = \frac{H_1(b)}{h_1(b)}$$

so the distribution of bids of the strong bidder H_1 dominates the distribution of bids of the weak bidder H_2 in terms of the reverse hazard rate. Thus, under the hypotheses of Proposition 4.4, H_1 also stochastically dominates H_2.

Why is it that the weak bidder bids more aggressively than does the strong bidder? To gain some intuition, it is useful to see why the opposite is impossible—that is, it cannot be that the strong bidder bids more aggressively than does the weak bidder. If for all x, $\beta_1(x) > \beta_2(x)$, then certainly the distribution H_1 of competing bids facing the weak bidder is stochastically higher than the distribution H_2 of competing bids facing the strong bidder. It is easy to see that all else being equal, a bidder who faces a stochastically higher distribution of bids—in the sense of reverse hazard rate dominance—will bid higher. It is also true that for a particular bidder, all else being equal, a higher realized value will lead to a higher bid. Now consider a particular bid b and suppose that $\beta_1(x_1) = \beta_2(x_2) = b$.

Since by assumption, the strong bidder bids more aggressively, it must be that the value at which the strong bidder bids b is lower than the value at which the weak bidder bids b—that is, $x_1 < x_2$. This means that, relative to the strong bidder, the weak bidder faces *both* a stochastically higher distribution of competing bids—H_1 versus H_2—and has a higher value—x_2 versus x_1. Since both forces cause bids to be higher, if it were optimal for the strong bidder to bid b when his value is x_1, it cannot be optimal for the weak bidder to bid b when his value is x_2. Thus, we have a contradiction.

Put another way, in equilibrium the two forces must balance each other. The weak bidder faces a stochastically higher distribution of competing bids than does the strong bidder, but the value at which any particular bid b is optimal for the weak bidder is lower than it is for the strong bidder.

ASYMMETRIC UNIFORM DISTRIBUTIONS

Equilibrium bidding strategies in asymmetric first-price auctions can be explicitly derived if the two value distributions are uniform but with differing supports. Specifically, suppose bidder 1's value X_1 is uniformly distributed on $[0, \omega_1]$ and bidder 2's value X_2 is uniformly distributed on $[0, \omega_2]$ and that $\omega_1 \geq \omega_2$. Then $F_i(x) = x/\omega_i$ and $f_i(x) = 1/\omega_i$, so the first-order condition (4.17) can be simplified as follows: for $i = 1, 2$ and $j \neq i$, for all $b \in (0, \overline{b})$,

$$\phi_i'(b) = \frac{\phi_i(b)}{\phi_j(b) - b} \tag{4.20}$$

which is equivalent to

$$\left(\phi_i'(b) - 1\right)\left(\phi_j(b) - b\right) = \phi_i(b) - \phi_j(b) + b$$

Adding the two equations for $i = 1, 2$ results in

$$\frac{d}{db}\left(\left(\phi_1(b) - b\right)\left(\phi_2(b) - b\right)\right) = 2b$$

and integrating this, we obtain

$$\left(\phi_1(b) - b\right)\left(\phi_2(b) - b\right) = b^2 \tag{4.21}$$

(The constant of integration is zero since $\phi_i(0) = 0$.) Since $\phi_i\left(\overline{b}\right) = \omega_i$,

$$\left(\omega_1 - \overline{b}\right)\left(\omega_2 - \overline{b}\right) = \overline{b}^2$$

so that

$$\overline{b} = \frac{\omega_1 \omega_2}{\omega_1 + \omega_2} \tag{4.22}$$

Now, using (4.21), the equations in (4.20) can be rewritten as follows: for $i=1,2$,

$$\phi_i'(b)=\frac{\phi_i(b)\,(\phi_i(b)-b)}{b^2} \tag{4.23}$$

with the advantage that in this form they are separable in the variables.

We now undertake a change of variables by defining $\xi_i(b)$ implicitly by

$$\phi_i(b)-b=\xi_i(b)b \tag{4.24}$$

so

$$\phi_i'(b)-1=\xi_i'(b)b+\xi_i(b)$$

With this substitution, the differential equation (4.23) becomes

$$\xi_i'(b)b+\xi_i(b)+1=\xi_i(b)\,(\xi_i(b)+1)$$

or

$$\frac{\xi_i'(b)}{\xi_i(b)^2-1}=\frac{1}{b}$$

the solution to which is easily verified to be

$$\xi_i(b)=\frac{1-k_ib^2}{1+k_ib^2},$$

where k_i is a constant of integration. Using (4.24) this becomes

$$\phi_i(b)=\frac{2b}{1+k_ib^2} \tag{4.25}$$

and since $\phi_i\left(\overline{b}\right)=\omega_i$, where $\overline{b}$ is defined in (4.22), we obtain that the constant of integration

$$k_i=\frac{1}{\omega_i^2}-\frac{1}{\omega_j^2} \tag{4.26}$$

The bidding strategies, obtained by inverting (4.25), are

$$\beta_i(x)=\frac{1}{k_ix}\left(1-\sqrt{1-k_ix^2}\right) \tag{4.27}$$

It is routine to verify that these form an equilibrium.[2] Figure 4.2 depicts the equilibrium bidding strategies when $\omega_1=\frac{4}{3}$ and $\omega_2=\frac{4}{5}$.

[2]Although we do not verify this here, it can be shown that this is the only equilibrium.

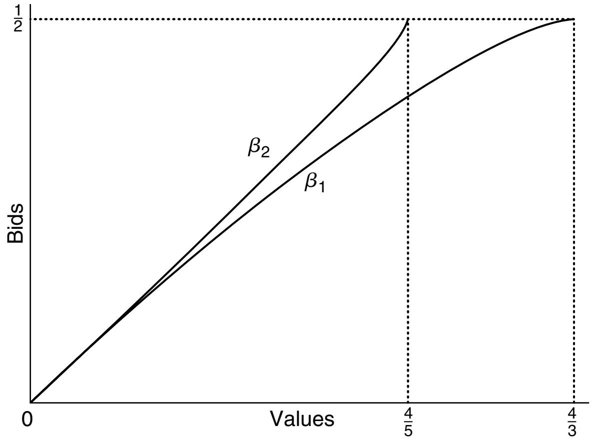

FIGURE 4.2 Equilibrium of an asymmetric first-price auction.

4.3.2 Revenue Comparison

We first use the equilibrium strategies derived previously to show that for some distributions, the revenue from a first-price auction may exceed that from a second-price auction.

Example 4.3. *With asymmetric bidders, the expected revenue in a first-price auction may exceed that in a second-price auction.*

As a special case of the example with values drawn from different uniform distributions, let $\alpha \in [0,1)$ and suppose that bidder 1's value X_1 is uniformly distributed on the interval $[0, 1/(1-\alpha)]$, whereas bidder 2's value X_2 is uniformly distributed on the interval $[0, 1/(1+\alpha)]$. (In the example depicted in Figure 4.2, $\alpha = 1/4$.)

We will compare the expected revenues accruing from a second-price auction to those from a first-price auction when $\alpha > 0$. Notice that when $\alpha = 0$, the situation is symmetric and the two auctions yield the same expected revenue.

REVENUE IN THE SECOND-PRICE AUCTION

It is a dominant strategy to bid one's value in a second-price auction and thus the distribution of the selling price in a second-price auction is

$$L_\alpha^{\text{II}}(p) = \text{Prob}\left[\min\{X_1, X_2\} \le p\right],$$

where $p \in \left[0, \frac{1}{1+\alpha}\right]$. We have

$$\begin{aligned} L_\alpha^{\text{II}}(p) &= F_1(p) + F_2(p) - F_1(p)F_2(p) \\ &= (1-\alpha)p + (1+\alpha)p - (1-\alpha)(1+\alpha)p^2 \\ &= 2p - \left(1-\alpha^2\right)p^2, \end{aligned}$$

which is *increasing* in α. Thus, in a second-price auction the expected selling price when $\alpha > 0$ is *lower* than the expected selling price when $\alpha = 0$.

REVENUE IN THE FIRST-PRICE AUCTION

Since $\omega_1 = 1/(1-\alpha)$ and $\omega_2 = 1/(1+\alpha)$, from (4.22) it follows that the highest amount that either bidder bids is $\overline{b} = \frac{1}{2}$. Moreover, the constants of integration in (4.26) are $k_1 = -4\alpha$ and $k_2 = 4\alpha$. Using (4.25) the inverse bidding strategies in equilibrium are: for all $b \in \left[0, \frac{1}{2}\right]$,

$$\phi_1(b) = \frac{2b}{1 - 4\alpha b^2}$$
$$\phi_2(b) = \frac{2b}{1 + 4\alpha b^2}$$

The distribution of the equilibrium prices in a first-price auction is thus

$$L_\alpha^{\mathrm{I}}(p) = \mathrm{Prob}\left[\max\left\{\beta_1(X_1), \beta_2(X_2)\right\} \le p\right],$$

where $p \in \left[0, \frac{1}{2}\right]$. We have

$$\begin{aligned} L_\alpha^{\mathrm{I}}(p) &= F_1(\phi_1(p)) \times F_2(\phi_2(p)) \\ &= (1-\alpha)\frac{2p}{1-4\alpha p^2} \times (1+\alpha)\frac{2p}{1+4\alpha p^2} \\ &= \frac{\left(1-\alpha^2\right)(2p)^2}{1-\alpha^2(2p)^4}, \end{aligned}$$

which is *decreasing* in α. Thus, in a first-price auction the expected selling price when $\alpha > 0$ is *higher* than the expected selling price when $\alpha = 0$.

For $\alpha = 0$, the expected selling price in the two auctions is the same. An increase in α leads to a decrease in the expected price in the second-price auction and an increase in the expected price in a first-price auction.

We have thus shown that in this example, for all $\alpha > 0$, the expected selling price in a first-price auction is *greater* than that in a second-price auction. (More generally, it can be shown that with asymmetric uniformly distributed values, the first-price auction is revenue superior for all ω_1 and ω_2.) ▲

A second example shows that the opposite ranking is also possible.

Example 4.4. *With asymmetric bidders, the expected revenue in a second-price auction may exceed that in a first-price auction.*

Suppose that bidder 1's value X_1 is distributed according to $F_1(x) = x - 1$ over $[1, 2]$ and bidder 2's value is distributed according to $F_2(x) = \exp\left(\frac{1}{2}x - 1\right)$

over $[0,2]$. (Note that we are departing from our assumption that lowest possible value for each bidder is 0. Also, F_2 has a mass point at 0.)

It may be verified that the bidding strategies $\beta_1(x) = x - 1$ and $\beta_2(x) = \frac{1}{2}x$ for the two bidders, respectively, constitute an equilibrium of the first-price auction. The bids range between 0 and 1 and it is routine to see that the distribution of prices in a first-price auction is

$$L^{\mathrm{I}}(p) = p\exp(p-1)$$

The expected revenue in a first-price auction $E\left[R^{\mathrm{I}}\right] \simeq 0.632$.

On the other hand, in a second-price auction the selling prices are distributed according to

$$L^{\mathrm{II}}(p) = \begin{cases} \exp\left(\frac{1}{2}p - 1\right) & p \leq 1 \\ (p-1) + (2-p)\exp\left(\frac{1}{2}p - 1\right) & p > 1 \end{cases}$$

over $[0,2]$. The expected revenue in a second-price auction $E\left[R^{\mathrm{II}}\right] \simeq 0.662$.

Thus, in this example, the expected selling price in a first-price auction is *less* than that in a second-price auction. ▲

4.3.3 Efficiency Comparison

As noted earlier, it is a weakly dominant strategy for a bidder to bid his or her value in a second-price auction—recall that this is true even when bidders are asymmetric—so the winning bidder is also the one with the highest value. Thus, the second-price auction is always *ex post* efficient under the assumption of private values.

In contrast, asymmetries inevitably lead to inefficient allocations in a first-price auction. Suppose that there are two bidders and (β_1, β_2) is an equilibrium of the first-price auction such that both strategies are continuous and increasing. Because the bidders are asymmetric—their values are drawn from different distributions—it will be the case that $\beta_1 \neq \beta_2$. Without loss of generality, suppose that $\beta_1(x) < \beta_2(x)$, and since both the strategies are continuous, for small enough $\varepsilon > 0$, it is also the case that $\beta_1(x+\varepsilon) < \beta_2(x-\varepsilon)$. This, of course, means that with positive probability the allocation is inefficient, since bidder 2 would win the auction even though he has a lower value than does bidder 1.

For future reference we record these observations as follows:

Proposition 4.5. *With asymmetrically distributed private values, a second-price auction always allocates the object efficiently, whereas with positive probability, a first-price auction does not.*

4.4 RESALE AND EFFICIENCY

In the previous section we saw that asymmetries among bidders lead to inefficient allocations in first-price auctions—with positive probability the winner of the auction is not the person who values the object the most. Achieving an efficient allocation may well be an important, or even primary, policy goal of the seller, especially if the seller is a government undertaking the privatization of some public asset. This seems to imply that such a seller should use an efficient auction—with private values, say a second-price auction—even if, as we have seen, it may bring lower revenues than an inefficient one, say a first-price auction. An argument against this point of view, in the Chicago school vein, is that even if the outcome of the auction is inefficient, postauction transactions among buyers—resale—will result in an efficient final allocation. Absent any transaction costs, the asset will be transferred into the hands of the person who values it the most. The conclusion is that the choice of the auction form is irrelevant to the efficiency question and one may as well select the auction format on other grounds—say, revenue. In this section we ask whether this is indeed the case. Does resale automatically lead to efficiency?

To examine the resale question in the simplest possible setting, consider the basic setup of the previous section. There are two bidders with values X_1 and X_2, which are independently distributed according to the functions F_1 and F_2, respectively, and for notational ease suppose that these have a common support $[0, \omega]$. The bidders are asymmetric, so $F_1 \neq F_2$. Suppose in addition that $E[X_1] \neq E[X_2]$, a condition that will hold generically.

In this context, let us first put forward the argument that a first-price auction followed by resale will lead to efficiency. Suppose β_1 and β_2 are equilibrium bidding strategies in the first-price auction and we know that these are increasing. Further, suppose that at the conclusion of the auction, both the bids—winning and losing—are publicly announced. This means that if the bids were b_1 and b_2, then at the conclusion of the auction, it would be commonly known that the buyers have values $x_1 = \beta_1^{-1}(b_1)$ and $x_2 = \beta_2^{-1}(b_2)$, respectively. If $b_1 > b_2$ but $x_1 < x_2$, the outcome of the auction would be inefficient, but since the values would be commonly known, so would the fact that there are some unrealized gains from trade. In particular, knowing that $x_1 < x_2$, bidder 1 could offer to resell the object to bidder 2 at some price between the two values. This would mean that ultimately the object ends up in the right hands.

This line of reasoning seems so simple as to be beyond question. It fails, however, to take into account that rational buyers will behave differently during the auction once they correctly foresee future resale possibilities. Let us see why.

To model the situation outlined above carefully, we need to be more specific about how resale actually takes place—that is, how the buyer and the seller settle on a price. Suppose that after learning what the losing bid was, the winner of the auction—and the new owner of the object—may, if he so wishes, resell the object to the other bidder by making a one-time take-it-or-leave-it offer. Notice that all bargaining power in this transaction resides with the (new) seller—that is, the winner of the auction. In particular, if bidder 1 wins the auction and subsequently learns that $x_1 < x_2$, then he can offer to sell the object to bidder 2 at a price $p = x_2$ (or just below), and this offer will be accepted. Bidder 1 will then make a profit of $x_2 - b_1$, whereas bidder 2's profit will be 0 (or just above). Of course, this is one of many possible ways in which the price may be determined. It is particularly simple and has the virtue that it ensures efficiency if the values are commonly known, so in some sense, it makes the best case for resale since there are no underlying transaction costs or delays. As an alternative, one could consider a situation in which the buyer makes a take-it-or-leave-it offer to the seller without affecting any of what follows.

Our main finding is that there cannot be an equilibrium of a first-price auction with resale in which the outcome of the auction completely reveals the values. Thus, the prospects of postauction resale make an efficient allocation impossible.

Suppose to the contrary that the first-price auction with resale has an efficient equilibrium. The equilibrium must specify both how bidders bid in the auction and what they do in the post-auction resale stage. Let β_1 and β_2 denote the bidding strategies in the first-price auction, and suppose that these are *increasing* with inverses $\phi_1 \equiv \beta_1^{-1}$ and $\phi_2 \equiv \beta_2^{-1}$. In the resale stage, if the announced bids b_1 and b_2 are such that $b_i > b_j$ but $x_i < \phi_j(b_j)$, then i makes a take-it-or-leave-it offer to sell the object to j at a price of $\phi_j(b_j)$. This is accepted if and only if $x_j \geq \phi_j(b_j)$. Otherwise, no offer is made.

The assumption that β_1 and β_2 are invertible means that after the winning and losing bids are announced, the values will become commonly known. Thus, if there are any unrealized gains from trade, resale will take place and the object will be allocated efficiently.

As a first step, notice that, as in the previous section, the bidding strategies must agree both at the lower and the upper end of the support of values—that is, $\beta_1(0) = 0 = \beta_2(0)$ and for some $\overline{b}$,

$$\beta_1(\omega) = \overline{b} = \beta_2(\omega) \tag{4.28}$$

Now suppose bidder 2 behaves according to the prescribed equilibrium strategy. Suppose bidder 1 has value x_1 but behaves as if his value were z_1—that is, he bids an amount $\beta_1(z_1)$ in equilibrium and in the resale stage also follows the given equilibrium strategy as if his value were z_1. Bidder 1's overall expected payment in equilibrium when he behaves as if his value is z_1 is[3]

$$\begin{aligned} m_1(z_1) = {} & F_2(\phi_2\beta_1(z_1))\,\beta_1(z_1) \\ & - \int_{z_1}^{\phi_2\beta_1(z_1)} \max\{z_1, x_2\} f_2(x_2)\, dx_2 \end{aligned} \tag{4.29}$$

[3]We write $\phi_2\beta_1(z_1)$ to denote $\phi_2(\beta_1(z_1))$.

The first term is bidder 1's expected payment to the seller. The second term is the result of monetary transfers between the bidders in the event that resale takes place. If $z_1 < X_2 < \phi_2\beta_1(z_1)$, then bidder 1 wins the auction and resells the object to bidder 2 at a price of $X_2 = \max\{z_1, X_2\}$. On the other hand, if $\phi_2\beta_1(z_1) < X_2 < z_1$, then bidder 1 loses the auction but purchases the object from bidder 2 at a price of $z_1 = \max\{z_1, X_2\}$.

The final allocation is efficient, so if bidder 1 behaves as if his value is z_1, the probability that he will get the object is just $F_2(z_1)$. Following the reasoning in the proof of the revenue equivalence principle (Proposition 3.1 on page 28), the expected payoff to bidder 1 from behaving as if his value is z_1 when it is actually x_1 is

$$F_2(z_1)\,x_1 - m_1(z_1)$$

In equilibrium, it is not optimal for bidder 1 to deviate, so we must have

$$F_2(x_1)\,x_1 - m_1(x_1) \geq F_2(z_1)\,x_1 - m_1(z_1)$$

The first-order condition for this optimization results in the differential equation

$$m_1'(x_1) = x_1 f_2(x_1)$$

and since $m_1(0) = 0$, it follows that

$$m_1(x_1) = \int_0^{x_1} x_2 f_2(x_2)\,dx_2 \tag{4.30}$$

Setting $z_1 = x_1$ in (4.29) and equating this with (4.30), we obtain that a necessary condition for a first-price auction followed by resale to be efficient is that for all x_1,

$$\begin{aligned} F_2(\phi_2\beta_1(x_1))\beta_1(x_1) - \int_{x_1}^{\phi_2\beta_1(x_1)} \max\{x_1, x_2\} f_2(x_2)\,dx_2 \\ = \int_0^{x_1} x_2 f_2(x_2)\,dx_2 \end{aligned} \tag{4.31}$$

Equation (4.31) says that the expected payment of bidder 1 in an equilibrium of the first-price auction with resale that is efficient—the expression on the left-hand side—is the same as that in an efficient *second*-price auction—the expression on the right-hand side. Indeed, it is just a version of the revenue equivalence principle extended to the asymmetric case: since the equilibrium outcomes of a first-price auction with resale are the same as those of a second-price auction, the expected payments in the two must be the same.[4]

[4]A general revenue equivalence principle for asymmetric situations is derived in Chapter 5 and can also be used to deduce equation (4.31).

Now $\beta_1(\omega) = \overline{b}$, so $\phi_2\beta_1(\omega) = \omega$. Setting $x_1 = \omega$ in (4.31), we obtain

$$\overline{b} = E[X_2]$$

But interchanging the roles of bidder 1 and 2, we similarly obtain that

$$\overline{b} = E[X_1]$$

and since $E[X_1] \neq E[X_2]$, this is a contradiction. We have thus argued that in an asymmetric first-price auction followed by resale, the bidding strategies β_1 and β_2 *cannot* be increasing everywhere. Bidders' equilibrium behavior in the auction cannot reveal their values completely, so resale transactions must take place under incomplete information.

From here it is a short step to see that because of this incomplete information, reaching an efficient allocation in all circumstances is impossible. In the interests of space, we only sketch the argument. Suppose the equilibrium strategies β_1 and β_2 are continuous but only nondecreasing (we have already ruled out the possibility that they are increasing). In other words, the strategies involve some "pooling"—that is, for at least one of the bidders—say, bidder 2—there is an interval of values $[x_2', x_2'']$ such that for all $x_2 \in [x_2', x_2'']$, the equilibrium bid $\beta_2(x_2)$ is the same, say b_2. We will now argue that pooling is also incompatible with efficient resale.

We first claim that there exists an $x_1 \in (x_2', x_2'')$ such that $\beta_1(x_1) \geq b_2$. Otherwise, for all x_1 such that $\beta_1(x_1) \geq b_2$, $x_1 \geq x_2''$. Let x_1'' be the smallest value for bidder 1 such that $\beta_1(x_1) = b_2$, and we know that $x_1'' \geq x_2''$. Now we must have $b_2 < x_1''$, since otherwise bidder 2 with value x_2'' would never bid b_2. The most that she can gain from winning the auction is $\max\{x_1'', x_2''\} = x_1''$, so she would never bid more than this amount. Since $b_2 < x_1''$, for ε small enough, $b_2 < x_1'' - \varepsilon$. But now notice that by bidding $b_2 + \varepsilon$ instead of below b_2, bidder 1 with value $X_1 = x_1'' - \varepsilon$ would gain a discrete amount—winning against all values X_2 in $[x_2', x_2'']$ to which he was losing previously—at an infinitesimal cost. This is a profitable deviation, so we have reached a contradiction.

Thus, there exists an $x_1 \in (x_2', x_2'')$ such that $\beta_1(x_1) \geq b_2$. If the realized values are x_1 and some $x_2 \in (x_1, x_2'']$, then bidder 1 wins the auction, but $x_1 < x_2$. The announcement of b_2 reveals to bidder 1 only that $X_2 \in [x_2', x_2'']$. He would then make a take-it-or-leave-it offer p that maximizes his expected profit. But such an offer would of necessity satisfy $p > x_1$ and so would be rejected with positive probability even though efficiency dictates that the object be transferred from bidder 1 to bidder 2. Thus, any equilibrium with pooling is inefficient.

If the resale price is determined instead by the buyer making a take-it-or-leave-it offer, the argument is virtually the same as the preceding one, except that the resale price is $\min\{X_1, X_2\}$ instead of $\max\{X_1, X_2\}$. All the other steps are identical.

We have thus argued the following:

Proposition 4.6. *With asymmetric bidders, a first-price auction followed by resale (at a take-it-or-leave-it price offered by one of the parties) does not result in efficiency.*

Proposition 4.6 casts doubt on the argument that resale will inevitably lead to efficiency. A seller whose goal is to ensure that the object ends up in the hands of the person who values it the most cannot rely on the "market" to do the job. The appropriate choice of an auction format remains very relevant—in order to assure efficiency, it is best to use an efficient auction.

PROBLEMS

4.1. (Risk-averse bidders) There are two bidders with private values which are distributed independently according to the uniform distribution $F(x)=x$ over $[0,1]$. Both bidders are risk-averse and have utility functions $u(z)=\sqrt{z}$. Find symmetric equilibrium bidding strategies in a first-price auction.

4.2. (Increase in risk aversion) Consider an N-bidder first-price auction where each bidder's value is distributed according to F. All bidders are risk averse with a utility function u that satisfies $u(0)=0$, $u'>0$, $u''<0$. Show that if one changed the utility function of all bidders from $u(z)$ to $\phi(u(z))$, where ϕ is an increasing and concave function satisfying $\phi(0)=0$, then this would lead to a higher symmetric equilibrium bidding strategy.

4.3. (Asymmetric first-price auction) Suppose that bidder 1's value X_1 is distributed according to $F_1(x)=\frac{1}{4}(x-1)^2$ over $[1,3]$ and bidder 2's value is distributed according to $F_2(x)=\exp\left(\frac{2}{3}x-2\right)$ over $[0,3]$.

a. Show that $\beta_1(x)=x-1$ and $\beta_2(x)=\frac{2}{3}x$ constitute equilibrium bidding strategies in a first-price auction.

b. Compare the expected revenues in a first- and second-price auction.

4.4. (Equilibrium with reserve price) Suppose that bidder 1's value X_1 is distributed uniformly on $[0,2]$ while bidder 2's value X_2 is distributed uniformly on $\left[\frac{3}{2},\frac{5}{2}\right]$. The object is sold via a first-price auction with a reserve price $r=1$. Verify that $\beta_1(x)=\frac{x}{2}+\frac{1}{2}$ and $\beta_2(x)=\frac{x}{2}+\frac{1}{4}$ constitute equilibrium strategies.

4.5. (Discrete values) Suppose that there is no uncertainty about bidder 1's value and $X_1=2$ always. Bidder 2's value, X_2, is equally likely to be 0 or 2.

a. Find equilibrium bidding strategies in a first-price auction. (Note that since values are discrete, the equilibrium will be in mixed strategies.)

b. Compare the revenues in a first- and second-price auction.

CHAPTER NOTES

The result that with symmetric risk-averse bidders the revenue from the first-price auction is greater than that in a second-price auction is due to Holt (1980), who studied the analogous problem in the context of bidding for procurement contracts. The key to Proposition 4.1, of course, is the fact that in a first-price auction, risk aversion causes bidders to increase their bids. The result that CARA bidders are indifferent between the first- and second-price auctions is due to Matthews (1987).

The material on auctions with budget constraints is based on Che and Gale (1998). This paper provides a sufficient condition on the primitives that guarantees that the first-price auction has an equilibrium of the sort hypothesized in Proposition 4.3. It also considers situations where the financial constraints need not be in the form of an absolute budget but may be somewhat more flexible. For instance, it may be that bidders face borrowing constraints so that they can borrow larger amounts only at higher marginal costs.

Vickrey (1961) himself pointed out that asymmetries among bidders may lead to inefficient allocations in a first-price auction. He studied a two-bidder asymmetric example in which bidder 1's value, a, was commonly known and bidder 2's value was uniformly distributed on $[0, 1]$. In this case, an equilibrium of the first-price auction involves randomization on the part of bidder 1, and Vickrey (1961) showed that depending on the value of a, the revenue from a first-price auction could be better or worse than that from a second-price auction. The derivation of equilibrium bidding strategies in first-price auctions with asymmetric uniformly distributed values is due to Griesmer, Levitan, and Shubik (1967). Plum (1992) has extended this to the more general class of asymmetric "power" distributions. Cheng (2006) identifies another family of distributions for which equilibrium strategies in an asymmetric first-price auction can be explicitly derived. It can be shown that an equilibrium exists in general under weak conditions on the distributions of values. Appendix G outlines some results in this direction. With a view to applications, Marshall, Meurer, Richard, and Stromquist (1994) have developed numerical techniques to compute bidding strategies in asymmetric auctions.

Maskin and Riley (2000a) have studied the equilibrium properties of asymmetric first-price auctions in more detail and have derived some sufficient conditions for one or the other auction to be revenue superior. Example 4.3 is taken from their paper, while Example 4.4 is due to Cheng (private communication, 2007).

Gupta and Lebrun (1999) have also studied a model of a first-price auction with resale but reach very different conclusions from those reached here. In particular, they do not find that resale will lead to inefficiency. This discrepancy is easily accounted for. Gupta and Lebrun (1999) assume that regardless of the outcome, at the conclusion of the auction the *values* of the bidders are publicly announced. Since after the auction is over, bidders are completely informed of each other's values, resale is inevitably efficient. The auctioneer, however, has no direct knowledge of bidders' values, so it is not clear how this is to be implemented. In contrast, in our model, we assumed that only the *bids* were announced and showed that if bidders take this into account, the bids would not reveal the values completely.

Garratt and Tröger (2006) study the possibility that a bidder with a known value of 0 may speculatively bid in a first-price auction when resale is present. Hafalir and Krishna (2008) construct equilibria of first-price auctions with resale and show that revenue from a first-price auction with resale exceeds that from a second-price auction. This result holds for all asymmetric distributions which satisfy a regularity condition (see Chapter 5 for a definition).

Haile (2003) studies a model of resale in which, at the time of bidding, buyers have only partial information regarding the true value. After the auction is over they receive a further signal that determines their actual values. The discrepancy between the estimated values at the time of the auction and the true values realized after the auction creates a motive for resale.