

COPYRIGHT NOTICE:

Paul Klemperer: Auctions: Theory and Practice

is published by Princeton University Press and copyrighted, © 2004, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

For COURSE PACK and other PERMISSIONS, refer to entry on previous page. For more information, send e-mail to permissions@pupress.princeton.edu

A Survey of Auction Theory*

*This chapter provides an elementary, non-technical survey of auction theory, by introducing and describing some of the critical papers in the subject. (The most important of these are reproduced in a companion book, Klemperer, Paul (ed.) (2000a) *The Economic Theory of Auctions*. Cheltenham, UK: Edward Elgar.) We begin with the most fundamental concepts, and then introduce the basic analysis of optimal auctions, the revenue equivalence theorem, and marginal revenues. Subsequent sections address risk aversion, affiliation, asymmetries, entry, collusion, multi-unit auctions, double auctions, royalties, incentive contracts, and other topics. Appendices contain technical details, some simple worked examples, and bibliographies. An Afterword to bring the survey up to date, and Exercises, are at the end of the chapter.¹*

1.1 INTRODUCTION

Auction theory is important for practical, empirical, and theoretical reasons.

First, a huge volume of goods and services, property, and financial instruments, are sold through auctions, and many new auction markets are being designed, including, for example, for mobile-phone licenses, electricity, and pollution permits.² Parts C and D of this volume discuss auction design in practice.

Second, auctions provide a very valuable testing-ground for economic theory—especially of game theory with incomplete information—which has been increasingly exploited in recent years. Major empirical research efforts

* This chapter (without the Afterword and Exercises) was originally published under the title *Auction Theory: A Guide to the Literature*, in the *Journal of Economic Surveys* 1999, 13, 227–286. (It is also reprinted in Dahiya, S. (ed.) (1999) *The Current State of Economic Science*. India: Spellbound Publications.) I would like to thank my friends and colleagues, including Mark Armstrong, Chris Avery, Alan Beggs, Sushil Bikhchandani, Simon Board, Jeremy Bulow, Peter Cramton, Nils Henrik von der Fehr, Tim Harford, Jon Levin, Ulrike Malmendier, Flavio Menezes, Meg Meyer, Paul Milgrom, John Morgan, John Riley, Mark Satterthwaite, Stuart Sayer, Daniel SgROI, Margaret Stevens, John Thanassoulis, Chris Wallace, Lucy White, Bob Wilson, and a referee for helpful advice.

¹ Other detailed treatments of the theory are in Krishna (2002) and Menezes and Monteiro (in preparation). For a more advanced analysis, see Milgrom (2004).

² See part D, McAfee and McMillan (1994, 1996), Klemperer (1998), and Milgrom (2004) for discussion of mobile-phone license auctions; Green and Newbery (1992), for example, discuss the use in the UK electricity market of the auction mechanism first analyzed in Klemperer and Meyer (1989); Klemperer et al. (forthcoming) discuss auctions for environmental improvements.

have focused on auctions for oil drilling rights, timber, and treasury bills,³ and there has also been an upsurge of interest in experimental work on auctions.⁴

Finally, auction theory has been the basis of much fundamental theoretical work: it has been important in developing our understanding of other methods of price formation, most prominently posted prices (as, e.g., observed in most retail stores) and negotiations in which both the buyer and seller are actively involved in determining the price. There are close connections between auctions and competitive markets.⁵ There is also a very close analogy between the theory of optimal auctions and the theory of monopoly pricing,⁶ and auction theory can also help develop models of oligopolistic pricing.⁷ Auction-theoretic models and techniques also apply to non-price means of allocation.⁸ The connections between auction theory and other parts of economic theory are the topic of part B of this volume.

1.1.1 Plan of This Chapter

This chapter provides an elementary survey of auction theory, by introducing and describing some of the critical papers in the subject. The most important of these are reproduced in a companion book, *The Economic Theory of Auctions*,⁹ for which this chapter was originally prepared.

For readers completely new to auction theory, the remainder of this section provides a brief resumé of the simplest concepts. The subsequent sections correspond to the sections into which *The Economic Theory of Auctions* is organized. Section 1.2 discusses the early literature, and section 1.3 introduces the more recent literature. Section 1.4 introduces the analysis of optimal auctions and auction theory's most fundamental result: the revenue equivalence theorem. It also describes how the concept of "marginal revenue" can inform auction theory. (Technical details are given in appendices.) Section 1.4 focuses on auction theory's basic model of a fixed set of symmetric, risk-neutral bidders with independent information who bid independently for a single object. Most of the remainder of this chapter is about the effects of relaxing one or more of these assumptions. Section 1.5 permits risk-aversion; section 1.6 allows for correlation or affiliation of bidders' information (with

³ See Laffont (1997).

⁴ See Kagel (1995).

⁵ See section 1.8.3.

⁶ See sections 1.4 and 2.4.

⁷ Appendix B of Bulow and Klemperer (1998) provides one illustration. See also section 2.5.

⁸ Queues and lobbying contests are examples of all-pay auction models; see, for example, section 2.2.3, Holt and Sherman (1982), and Riley (1989b). The war of attrition can also be modeled as a kind of all-pay auction; see sections 1.13.4 and 2.2.2, and Bulow and Klemperer (1999). Insights from auction theory can explain rationing; see sections 1.8.1 and 2.3.2, and Gilbert and Klemperer (2000).

⁹ Klemperer (2000a).

technical details in an appendix); section 1.7 analyzes cases with asymmetric bidders; section 1.8 considers bidders who have costs of entering an auction, and addresses other issues pertaining to the number of bidders; section 1.9 asks what is known if there are possibilities for collusion among bidders; and section 1.10 considers multi-unit auctions. Section 1.11 looks at auctions for incentive contracts, and auctions in which contestants may bid on royalty rates or quality levels in addition to prices. Section 1.12 reviews the literature on double auctions, and section 1.13 briefly considers some other important topics including budget constraints, externalities between bidders, jump bidding, the war of attrition, and competing auctioneers. Section 1.14 is about testing the theory, and section 1.15 concludes. Appendices 1.A, 1.B, and 1.C provide technical details about the revenue equivalence theorem, marginal revenues, and affiliation, respectively. Appendix 1.D provides some simple worked examples illustrating these appendices. Appendix 1.E provides a bibliography organized according to the sections of this chapter. An Afterword to bring the survey up to date, and Exercises, are at the end of the chapter.

1.1.2 The Standard Auction Types

Four basic types of auctions are widely used and analyzed: the ascending-bid auction (also called the open, oral, or English auction), the descending-bid auction (used in the sale of flowers in the Netherlands and so also called the Dutch auction by economists), the first-price sealed-bid auction, and the second-price sealed-bid auction (also called the Vickrey auction by economists).^{10,11} In describing their rules we will focus for simplicity on the sale of a single object.

In the ascending auction, the price is successively raised until only one bidder remains, and that bidder wins the object at the final price. This auction can be run by having the seller announce prices, or by having the bidders call out prices themselves, or by having bids submitted electronically with the best current bid posted. In the model most commonly used by auction theorists (often called the Japanese auction), the price rises continuously while bidders gradually quit the auction. Bidders observe when their competitors quit, and once someone quits, she is not let back in. There is no possibility for one bidder to preempt the process by making a large “jump bid”. We will assume this model of the ascending auction except where stated otherwise.¹²

¹⁰ Confusingly, the second-price sealed-bid auction is sometimes called a Dutch auction by investment bankers.

¹¹ Cassady’s (1967) book provides a detailed, although now somewhat dated, account of many of the contemporaneous auction institutions.

¹² Antiques and artwork are commonly sold using versions of the ascending auction, and houses are sometimes sold this way, too. Bikhchandani and Riley (1991) discuss different types of ascending auction. See also section 1.13.3.

The descending auction works in exactly the opposite way: the auctioneer starts at a very high price, and then lowers the price continuously. The first bidder who calls out that she will accept the current price wins the object at that price.¹³

In the first-price sealed-bid auction each bidder independently submits a single bid, without seeing others' bids, and the object is sold to the bidder who makes the highest bid. The winner pays her bid (i.e., the price is the highest or "first" price bid).¹⁴

In the second-price sealed-bid auction, also, each bidder independently submits a single bid, without seeing others' bids, and the object is sold to the bidder who makes the highest bid. However, the price she pays is the *second*-highest bidder's bid, or "second price". This auction is sometimes called a Vickrey auction after William Vickrey, who wrote the seminal (1961) paper on auctions.¹⁵

For reasons we will explain shortly, the ascending and descending auctions are sometimes referred to as open second-price and open first-price auctions, respectively.

1.1.3 The Basic Models of Auctions

A key feature of auctions is the presence of asymmetric information.¹⁶ (With perfect information most auction models are relatively easy to solve.)

¹³ For example, in Dutch flower auctions, the potential buyers all sit in a room at desks with buzzers connected to an electronic clock at the front of the room. The interior of the clock has information about what is being sold and the price at which the auction starts. Once the auction begins, a series of lights around the edge of the clock indicate to what percentage of the original asking price the good has fallen. As soon as one bidder buzzes in, she gets the flowers at the price indicated on the clock. (Except that, if there are several lots of the same flowers from the same seller available that morning, the buyer can choose to buy only some of the available lots, and the rest will be re-auctioned.) Fish are sold in a similar way in Israel, as is tobacco in Canada.

¹⁴ First-price sealed-bid auctions are used in auctioning mineral rights in government-owned land; they are also sometimes used in the sales of artwork and real estate. This method is also often used in procurement (i.e., competing contractors submit prices and the *lowest* bidder wins and receives her price for fulfilling the contract). UK Treasury securities are sold through the multi-unit equivalent of the first-price auction (every winner pays her own bid), and US Treasury auctions used to be run this way too, though recently the US Treasury has also been using a multi-unit version of the second-price sealed-bid auction.

¹⁵ This auction form is used for most auctions of stamps by mail, and is also used for other goods in some auctions on the internet (see Lucking-Reiley, 2000), but it is much less commonly used than the other standard forms (see Rothkopf, Teisberg, and Kahn, 1990 for some discussion why); it is commonly studied in part because of its attractive theoretical properties. A multi-unit version is sometimes used by governments when selling foreign exchange and by companies when buying back shares. Economists usually model the multi-unit version by assuming the price paid is the highest losing bid, since this has theoretical properties analogous to those of the single-unit second-price case. In practice the price paid is often that of the lowest winning bidder.

¹⁶ The appropriate concept of equilibrium is therefore Bayesian–Nash equilibrium. That is, each player's strategy is a function of her own information, and maximizes her expected payoff given other players' strategies and given her beliefs about other players' information. See, for example, Gibbons (1992).

In the basic *private-value* model each bidder knows how much she values the object(s) for sale, but her value is private information to herself.

In the *pure common-value* model, by contrast, the actual value is the same for everyone, but bidders have different private information about what that value actually is. For example, the value of an oil lease depends on how much oil is under the ground, and bidders may have access to different geological “signals” about that amount. In this case a bidder would change her estimate of the value if she learnt another bidder’s signal, in contrast to the private-value case in which her value would be unaffected by learning any other bidder’s preferences or information.

A general model encompassing both these as special cases assumes each bidder receives a private information signal, but allows each bidder’s value to be a general function of *all* the signals.¹⁷ For example, your value for a painting may depend mostly on your own private information (how much you like it) but also somewhat on others’ private information (how much they like it) because this affects the resale value and/or the prestige of owning it.

1.1.4 Bidding in the Standard Auctions

Consider first the descending auction. Note that although we described this as a dynamic game, each bidder’s problem is essentially static. Each bidder must choose a price at which she will call out, conditional on no other bidder having yet called out; and the bidder who chooses the highest price wins the object at the price she calls out. Thus this game is strategically equivalent to the first-price sealed-bid auction,¹⁸ and players’ bidding functions are therefore exactly the same.¹⁹ This is why the descending auction is sometimes referred to as an open first-price auction.

Now with private values, in the ascending auction, it is clearly a dominant strategy to stay in the bidding until the price reaches your value, that is, until you are just indifferent between winning and not winning. The next-to-last person will drop out when her value is reached, so the person with the highest

¹⁷ That is, bidder i receives signal t_i and would have value $v_i(t_1, \dots, t_n)$ if all bidders’ signals were available to her. In the private-value model $v_i(t_1, \dots, t_n)$ is a function only of t_i . In the pure common-value model $v_i(t_1, \dots, t_n) = v_j(t_1, \dots, t_n)$, for all t_1, \dots, t_n . (If i ’s actual value $V_i(t_1, \dots, t_n, s_1, \dots, s_k)$ is also a function of other information s_1, \dots, s_k , then of course $v_i(t_1, \dots, t_n) = E\{V_i(t_1, \dots, t_n, s_1, \dots, s_k) \mid t_1, \dots, t_n\}$ is just i ’s estimated value, but for most purposes it does not matter whether $v_i(t_1, \dots, t_n)$ is an estimated or an actual value.)

¹⁸ That is, the set of strategies available to a player is the same in the descending auction as in the first-price sealed-bid auction. Choosing any given bid yields the same payoffs in both games as a function of the other players’ bids.

¹⁹ To solve for the bidding strategies, the direct method is to examine a player’s first-order condition in which she trades off her probability of winning (which increases with her bid) with her profit conditional on winning (which decreases with her bid). Note 121 illustrates the method. For the independent-signal case a faster and more elegant approach is to use the revenue equivalence theorem, see Appendix 1.A. Appendix 1.D gives examples.

value will win at a price equal to the value of the second-highest bidder. Furthermore, a little reflection shows that in a second-price sealed-bid private-values auction it is optimal for a player to bid her true value, whatever other players do.²⁰ In other words “truth telling” is a dominant strategy equilibrium (and so also a Nash equilibrium), so here, too, the person with the highest value will win at a price equal to the value of the second-highest bidder. This is why the ascending auction is sometimes referred to as an open second-price auction. However, this equivalence applies only for private values, or if there are just two bidders. With any common components to valuations and more than two bidders, players learn about their values from when other players quit an ascending auction and condition their behavior on this information.

A key feature of bidding in auctions with common-values components is the *winner’s curse*: each bidder must recognize that she wins the object only when she has the highest signal (in symmetric equilibrium). Failure to take into account the bad news about others’ signals that comes with any victory can lead to the winner paying more, on average, than the prize is worth, and this is said to happen often in practice. In equilibrium, bidders must adjust their bids downwards accordingly.

Appendix 1.D provides examples of equilibrium bidding strategies (and the winner’s curse) in the standard auctions, in both private- and common-value contexts.

1.1.5 Terminology

Since the equivalence of descending and first-price sealed-bid auctions is completely general in single-unit auctions, and ascending and second-price sealed-bid auctions are also equivalent under many conditions (and have similar properties more broadly) we will often refer to the two kinds of auctions simply as *first-price* and *second-price*, respectively.

Also, we will refer to any model in which a bidder’s value depends to some extent on other bidders’ signals as a *common-value* model. However, note that some authors reserve the term “common-value” to refer only to the special case when all bidders’ actual values are identical functions of the signals (what we

²⁰ To confirm this, consider bidding $v - x$ when your true value is v . If the highest bid other than yours is w , then if $v - x > w$ you win the auction and pay w , just as if you bid v . If $w > v$ you lose the auction and get nothing, just as if you bid v . But if $v > w > v - x$, bidding $v - x$ causes you to lose the auction and get nothing, whereas if you had bid v , you would have won the auction and paid w for a net surplus of $v - w$. So you never gain, and might lose, if you bid $v - x$.

Now consider bidding $v + x$ when your true value is v . If the highest bid other than yours is w , then if $v > w$ you win and pay w , just as if you bid v . If $w > v + x$ you lose and pay nothing, just as if you bid v . But if $v + x > w > v$, having bid $v + x$ causes you to “win” an auction you otherwise would have lost, and you have to pay $w > v$ so you get negative surplus. So bidding $v + x$ may hurt you compared with bidding v , but it never helps you.

called the pure common-value case). Also (and somewhat inconsistently) we will use the term *almost common values* to mean almost pure common values.

Finally, there is no formal distinction between normal auctions, in which the auctioneer is the seller and the bidders are buyers who have values for the object(s) sold, and procurement auctions, where the auctioneer is a buyer and the bidders are sellers who have costs of supplying the object(s) bought. To avoid confusion we will generally adopt the former perspective (that the auctioneer is the seller) even when discussing papers that are couched in terms of the latter perspective.

1.2 EARLY LITERATURE

Auctions have been used from time immemorial,²¹ but they entered the economics literature relatively recently. Remarkably, the first treatment that recognized the game-theoretic aspects of the problem,²² Vickrey (1961), also made enormous progress in analyzing it including developing some special cases of the celebrated *Revenue Equivalence Theorem* (see below). Vickrey's 1961 and 1962 papers were deservedly a major factor in his 1996 Nobel prize,²³ and the 1961 paper, especially, is still essential reading.

Other influential early work was performed by Shubik and his co-authors, and by Wilson and his student, Ortega Reichert.

Griesmer, Levitan, and Shubik (1967) analyze the equilibrium of a first-price auction in which contestants' valuations are drawn from uniform distributions with different supports, while Wilson (1969) introduced the (pure) common-value model and developed the first closed-form equilibrium analysis of the winner's curse.²⁴

²¹ Shubik (1983) provides an attractively written historical sketch going back to the Babylonian and Roman empires. Most famously, the whole Roman empire was sold by ascending auction in A.D. 193 by the Praetorian Guards; the winner, and therefore next Emperor, was Didius Julianus who reigned for just over two months before being overthrown and executed by Septimius Severus (an early and sad case of the winner's curse); see also notes 56, 106, and 108, and see Gibbon (1776), volume I, chapter V for an account.

²² There are slightly earlier studies in the operations research literature, especially Friedman (1956), but these treat the problem decision-theoretically with bidders estimating opponents' bidding strategies based on a naive model of past behavior.

²³ He shared the prize with Jim Mirrlees whose 1971 paper, although couched in the context of the theory of optimal income taxation, developed techniques that were to prove critical to the later analysis of auctions. Vickrey (1976) makes additional contributions to auction theory, including sketching the "simultaneous ascending auction" later proposed by McAfee, Milgrom, and Wilson for the recent FCC auctions of radio spectrum licenses (see note 78).

²⁴ Rothkopf (1969) addresses a similar problem to Wilson's, but implicitly restricts bidders' strategies to multiples of their estimated values (see Rothkopf, 1980). Capen, Clapp, and Campbell (1971) is a well-known, more popular account of the winner's curse in practice that was immensely important in influencing bidding practice. Wilson (1967) is a precursor to Wilson (1969), but with a less natural equilibrium concept.

Ortega Reichert's (1968a) PhD thesis contains the seeds of much future work, but the time was perhaps not ripe for it, and it unfortunately never reached publication. It is a pleasure to be able to publish a small part of it in *The Economic Theory of Auctions*: the chapter we publish considers a sequence of two first-price auctions in which the correlation of players' values for the two objects means that a player's bid for the first object conveys information about her value for the second object, and hence about her likely second bid. (We also publish a short explanatory foreword to the chapter.²⁵) This analysis of a signaling game was enormously influential in, for example, guiding Milgrom and Roberts' (1982) analysis of limit pricing.²⁶

However, with the exception of Vickrey's first (1961) article, these are no longer papers for the beginner.

1.3 INTRODUCTION TO THE RECENT LITERATURE

The full flowering of auction theory came only at the end of the 1970s with critical contributions from Milgrom, in papers both on his own and with Weber; from Riley, in papers with Maskin and with Samuelson; and from Myerson, among others, in addition to more from Wilson. These and contemporaneous contributions rapidly moved the field close to its current frontier. A very readable introduction to the state of the field by the late 1980s is in McAfee and McMillan (1987a). Another helpful introductory article is Maskin and Riley (1985) which manages to convey many of the key ideas in a few pages by focusing on the case with just two possible types of each of just two bidders.^{27,28}

1.4 THE BASIC ANALYSIS OF OPTIMAL AUCTIONS, REVENUE EQUIVALENCE, AND MARGINAL REVENUES

Roughly simultaneously, Myerson (1981) and Riley and Samuelson (1981) showed that Vickrey's results about the equivalence in expected revenue of different auctions apply very generally:²⁹

²⁵ See also the brief exposition of this work in Section 3.1 of Weber (1983) (see section 1.10.3 of this survey, and reprinted in the corresponding part of *The Economic Theory of Auctions*).

²⁶ Personal communication from Paul Milgrom.

²⁷ A caveat is that the effects of correlated types cannot properly be discussed with just two types, and this section of the paper is a little flawed and confusing. However Riley (1989a) has a nice discussion of correlation with just *three* possible types of each bidder. See also Appendix 1.C.

²⁸ Other valuable survey material includes Milgrom (1985, 1987, 1989), Weber (1985), Riley (1989b), Maskin (1992), Wilson (1992), Harstad and Rothkopf (1994), Rothkopf (1994), and Wolfstetter (1996).

²⁹ Ortega Reichert (1968a) and Holt (1980) made some earlier extensions of Vickrey's work. Harris and Raviv (1981) covers much of the same ground as Myerson and Riley and Samuelson.

*Assume each of a given number of risk-neutral potential buyers of an object has a privately known signal independently drawn from a common, strictly increasing, atomless distribution. Then any auction mechanism in which (i) the object always goes to the buyer with the highest signal, and (ii) any bidder with the lowest-feasible signal expects zero surplus, yields the same expected revenue (and results in each bidder making the same expected payment as a function of her signal).*³⁰

Note that the result applies both to private-value models (in which a bidder's value depends only on her own signal), and to more general common-value models provided bidders' signals are independent.

Thus all the "standard" auctions, the ascending, the descending, the first-price sealed-bid, and the second-price sealed-bid, yield the same expected revenue under the stated conditions, as do many non-standard auctions such as an "all-pay" auction (in which every competitor pays her bid but only the highest bidder wins the object, as in a lobbying competition).³¹

This *Revenue Equivalence Theorem* result is so fundamental, so much of auction theory can be understood in terms of it, and at root the proof is so simple, that we offer an elementary derivation of it in Appendix 1.A. Any reader who is unfamiliar with the result, or who is under any misapprehension that it is a difficult one, is strongly urged to begin here.³²

Riley and Samuelson's proof is less direct than that of Appendix 1.A, but is still a simpler read than Myerson's, and Riley and Samuelson give more illustrations. However, Myerson offers the most general treatment, and also develops the mathematics used to prove revenue equivalence a little further to show how to derive optimal auctions (i.e., auctions that maximize the seller's expected revenue) for a wide class of problems (see below).

Although this work was a remarkable achievement, there seemed to be little relationship to traditional price theory, which made the subject a difficult one for many economists. Bulow and Roberts (1989) greatly simplified the analysis of optimal auctions by showing that the problem is, in their own words, "essentially equivalent to the analysis of standard monopoly third-degree price discrimination. The auctions problem can therefore be understood by applying the usual logic of marginal revenue versus marginal cost."

³⁰ This is not the most general statement. See Appendix 1.A. To see the necessity of a strictly increasing or atomless distribution, see note 117. See Riley (1989a) for revenue equivalence results for discrete distributions.

³¹ Other examples that can be modeled as all-pay auctions include queues (see Holt and Sherman (1982)), legal battles (see Baye, Kovenock, and De Vries (1997)), and markets with consumer switching costs in which firms compete for the prize of selling to new unattached consumers by lowering their prices to their old locked-in customers (see especially Appendix B of Bulow and Klemperer (1998) which explicitly uses the revenue equivalence theorem, and also Rosenthal (1980) and more generally, Klemperer (1995)). The war of attrition is also a kind of all-pay auction (see section 1.13.4) and Bulow and Klemperer (1999).

³² The appendix also gives an example of solving for bidding strategies in more complex auctions by using revenue equivalence with an ascending auction.

In particular, it is helpful to focus on bidders' "marginal revenues". Imagine a firm whose demand curve is constructed from an arbitrarily large number of bidders whose values are independently drawn from a bidder's value distribution. When bidders have independent private values, a bidder's "marginal revenue" is defined as the marginal revenue of this firm at the price that equals the bidder's actual value. Bulow and Roberts follow Myerson to show that under the assumptions of the revenue equivalence theorem *the expected revenue from an auction equals the expected marginal revenue of the winning bidder(s)*.

Bulow and Klemperer (1996)³³ provide a simpler derivation of this result that also generalizes its application.³⁴ We give an elementary exposition of this material in Appendix 1.B.

So in an optimal auction the objects are allocated to the bidders with the highest marginal revenues, just as a price-discriminating monopolist sells to the buyers with the highest marginal revenues (by equalizing the lowest marginal revenues sold to across different markets). And just as a monopolist should not sell below the price where marginal revenue equals marginal cost, so an auctioneer should not sell below a reserve price set equal to the value of the bidder whose marginal revenue equals the value to the auctioneer of retaining the unit. (The marginal revenue should be set equal to zero if the auctioneer, or monopolist, is simply maximizing expected revenues.)

These principles indicate how to run an optimal auction in the general case.³⁵ Furthermore, when bidders are symmetric (i.e., when their signals are drawn from a common distribution), any "standard" auction sells to the bidder with the highest signal. Therefore, if bidders with higher signals have higher marginal revenues—in the private-value context this is just equivalent to the usual assumption that a monopolist's marginal revenue is downward sloping³⁶—then the winning bidder has the highest marginal revenue. So under the assumptions of the revenue equivalence theorem, and if bidders with higher signals have higher marginal revenues, *all the standard auctions are optimal if the seller imposes the optimal reserve price*.

³³ Also discussed in section 1.8.2, and reprinted in the corresponding part of *The Economic Theory of Auctions*.

³⁴ Bulow and Klemperer show how the result extends with common values, non-independent private information, and risk-aversion, while Bulow and Roberts restrict attention to the risk-neutral, independent, private-value, framework. See Appendix 1.B.

The main thrust of Bulow and Klemperer's analysis is to develop a result about the value to an auctioneer of an additional bidder relative to the importance of constructing an optimal auction (see section 1.8.2).

³⁵ See Myerson (1981) and Bulow and Roberts (1989) for details.

³⁶ This amounts to the assumption that the monopolist's demand curve (or bidder's distribution function) is not too convex.

The assumption that bidders with higher signals have higher marginal revenues is more stringent in common-value contexts. See note 54.

Much of auction theory can be most easily understood by thinking in terms of marginal revenues and the relationship to the conditions for revenue equivalence; this chapter emphasizes this perspective.

1.5 RISK-AVERSION

It is easy to see how risk-aversion affects the revenue equivalence result: in a second-price (or an ascending) auction, risk-aversion has no effect on a bidder's optimal strategy which remains to bid (or bid up to) her actual value.³⁷ But in a first-price auction, a slight increase in a player's bid slightly increases her probability of winning at the cost of slightly reducing the value of winning, so would be desirable for a risk-averse bidder if the current bidding level were optimal for a risk-neutral bidder. So risk-aversion makes bidders bid more aggressively in first-price auctions. Therefore, since the standard auctions were revenue equivalent with risk-neutral bidders, a risk-neutral seller faced by risk-averse bidders prefers the first-price auction to second-price sealed-bid or ascending auctions.

What if the auctioneer is risk-averse but the buyers are risk-neutral? Observe that the winner pays a price set by the runner-up in a second-price or ascending auction and, by revenue equivalence, must bid the expectation of this price in a first-price auction. That is, conditional on the winner's actual information, the price is fixed in the first-price auction, and random but with the same mean in the second-price auction. So also unconditional on the winner's information, the price is riskier (but with the same mean) in the second-price auction. So a risk-averse seller prefers the first-price auction to the second-price sealed-bid auction and, for a similar reason, prefers the second-price sealed-bid auction to an ascending open auction.³⁸

In another of the crucially important papers of the early 1980s, Maskin and Riley (1984) develop and generalize these results and then go on to consider the design of optimal auctions when the seller is risk-neutral and the buyers are risk-averse.³⁹

However, although first-price auctions lead to higher prices with risk-averse buyers, this does not mean risk-averse buyers prefer second-price or ascending auctions since, as noted above, prices in the first-price auction are less risky. Matthews (1987) takes up the buyer's viewpoint; in fact, buyers are just indifferent with constant absolute risk aversion and tend to prefer the first-price auction if they have increasing absolute risk aversion or "affiliated

³⁷ In a sealed-bid second-price auction with common-value components, a bidder bids her expected utility conditional on being tied for highest bidder (see Appendices 1.C and 1.D).

³⁸ See Waehrer, Harstad, and Rothkopf (1998) for the fullest exposition of the case of a risk-averse auctioneer.

³⁹ Matthews (1983) has many similar results. Holt (1980) is an earlier treatment.

values” (see next section).⁴⁰ These results can be developed by generalizing the revenue equivalence result to a “utility equivalence” result that applies for risk-averse bidders.⁴¹

1.6 CORRELATION AND AFFILIATION

Another crucial assumption in the basic analysis of optimal auctions is that each bidder’s private information is independent of competitors’ private information. We now relax this assumption while reverting to the assumption of risk-neutrality.

Section 7 of Myerson’s extraordinary (1981) paper provides a very simple and instructive example showing that if bidders’ private information is correlated, then the seller can construct a mechanism that yields for herself the entire social surplus that would be feasible if bidders’ information were fully public! The mechanism offers each bidder a schedule of bets among which she is required to choose if she is to participate. For any given private information, the best of these bets will yield her exactly zero surplus in expectation, and by choosing it she is revealing her type so that her surplus can be fully and efficiently extracted. We give an example in Appendix 1.C.

Crémer and McLean (1985) show that Myerson’s result is very general, although it does seem to rely heavily on assumptions such as risk-neutrality of both the bidders and the seller, common knowledge of the distributions from which bidders’ signals are drawn, the inability of bidders to collude, and the ability of the seller to credibly and costlessly communicate and enforce the auction’s results (often including extracting large payments from losing bidders).⁴²

Since the “optimal mechanisms” seem unrealistic in this environment, how do standard auctions compare? Milgrom and Weber’s remarkable (1982a) paper addresses this question in the context of developing a general theory of auctions with affiliated information. (Very roughly, bidders’ signals are *affiliated* if a high value of one bidder’s signal makes high values of other bidders’ signals more likely.⁴³) Since this paper is both very important in the literature and quite challenging for many readers, we give an elementary exposition of the main results in Appendix 1.C, by starting from the revenue equivalence argument developed in Appendix 1.A.

⁴⁰ Matthews’ paper is also important for its analysis of the case where the number of buyers is unknown. See section 1.8.4.

⁴¹ See Robert, Laffont, and Loisel (1994).

⁴² Crémer and McLean (1988), McAfee, McMillan, and Reny (1989), and McAfee and Reny (1992) show the result in even greater generality. Esö (1999) argues the result may not be too sensitive to bidder risk-aversion.

⁴³ See Appendix 1.C for a precise definition.

The main results are that ascending auctions lead to higher expected prices than sealed-bid second-price auctions, which in turn lead to higher expected prices than first-price auctions.⁴⁴ The intuition is that the winning bidder's surplus is due to her private information. The more the price paid depends on others' information (the price depends on all other bidders' information in an ascending auction with common-value elements, and on one other bidder's information in a second-price sealed-bid auction), the more closely the price is related to the winner's information, since information is affiliated. So the lower is the winner's information rent and hence her expected surplus, and so the higher is the expected price.

For the same reason, if the seller has access to any private source of information, her optimal policy is to pre-commit to revealing it honestly. The general principle that expected revenue is raised by linking the winner's payment to information that is affiliated with the winner's information, is known as the Linkage Principle.⁴⁵

One of the more striking results of the basic analysis of optimal auctions is that if bidders have independent private values, the seller's reserve price is both independent of the number of bidders, and also well above the seller's cost. The reason is that the optimal reserve price is where marginal revenue equals the seller's cost, and a bidder's marginal revenue is independent of other bidders' marginal revenues when values are independent. However, if valuations are affiliated, more bidders implies more certainty about any one bidder's valuation conditional on other bidders' information, hence flatter marginal revenue curves, so a far higher proportion of bidders have marginal revenues in excess of the seller's cost.⁴⁶ So the reserve price must be set lower. Levin and Smith (1996a) show that the optimal reserve price converges to the seller's true value as the number of bidders grows.

1.7 ASYMMETRIES

Along with risk-neutrality, and independent private information, a third crucial assumption of the revenue equivalence theorem is that buyers' private

⁴⁴ Appendix 1.D gives an example.

⁴⁵ See also the discussion of Milgrom and Weber (2000) in section 1.10.3.

⁴⁶ More precisely, consider setting an optimal reserve price after seeing all but the highest valuation. Affiliation implies the highest value is likely to be close to the second-highest, so the demand curve formed by a continuum of bidders with valuations drawn from the highest bidder's value distribution conditional on all others' values is rather flat just above the second value. Thus the final part of the highest-bidder's marginal revenue curve, conditional on all that has been observed, is also rather flat around the second-highest valuation. So even if the reserve price could be set based on all this information, it would usually be set very low. Hence it will also be set very low if it must be set prior to the auction.

(Note that even with independent signals the reserve price should optimally be set after the auction if there are common-value components to valuations. See, for example, Bulow and Klemperer (1996) for how to set the optimal reserve price in this case.)

values or signals were drawn from a common distribution. We now discuss relaxing the symmetry assumption.⁴⁷

1.7.1 Private Value Differences

Myerson (1981) and Bulow and Roberts (1989) showed that a revenue-maximizing auction allocates objects to the bidder(s) with the highest marginal revenue(s) rather than to those with the highest value(s). Recall from the standard theory of demand that a buyer on a given demand curve has a higher marginal revenue than any buyer with the same valuation on a demand curve that is higher everywhere due to being shifted out by a fixed amount horizontally. So a revenue-maximizing auctioneer typically discriminates in favor of selling to bidders whose values are drawn from lower distributions, that is, “weaker” bidders. McAfee and McMillan (1989) develop this point in a procurement context.⁴⁸

Since in a first-price auction a bidder whose value is drawn from a weaker distribution bids more aggressively (closer to her actual value) than a bidder from a stronger distribution,⁴⁹ a first-price auction also discriminates in favor of selling to the weaker bidder, in contrast to a second-price (or ascending) auction which always sells to the bidder with the higher valuation (in a private-values model). So it is plausible that a first-price auction may be more profitable in expectation, even though less allocatively efficient, than a second-price auction, when all the assumptions for revenue equivalence except symmetry are satisfied.⁵⁰ This is in fact often, though not always, the case. The large variety of different possible kinds of asymmetries makes it difficult to develop general results, but Maskin and Riley (2000b) make large strides.⁵¹ A very useful brief discussion in Maskin and Riley (1985) summarizes the situation as “roughly speaking, the sealed-bid auction generates more revenue than the open [second-price] auction when bidders have distributions with the same shape (but different supports), whereas the open auction dominates when,

⁴⁷ For results about the existence and uniqueness of equilibria in first-price auctions see Lebrun (1996), Maskin and Riley (2000a, forthcoming), Athey (2001), and Lizzeri and Persico (2000) (who consider a broader class of games). Some similar results can be developed following Bulow, Huang, and Klemperer (1995). Wilson (1998) derives explicitly the equilibrium of an ascending auction for a model with both private- and common-value components which allows asymmetries.

⁴⁸ See also Rothkopf, Harstad, and Fu (2003).

⁴⁹ In a first-price auction the first-order condition of a bidder with value v considering raising her bid, b , by a small amount Δb that would raise her probability of winning, p , by a small amount Δp sets $(v - b)\Delta p - p\Delta b = 0$. Weaker bidders have smaller probabilities of winning, p , and hence smaller “profit margins”, $v - b$, when they do win.

⁵⁰ Maskin (1992) shows an ascending auction is efficient for a single good, even when valuations have common-value components, under a broad class of assumptions.

⁵¹ The earliest analyses of asymmetric cases are in Vickrey (1961) and Griesmer, Levitan, and Shubik (1967). Marshall, Meurer, Richard, and Stromquist (1994), and Riley and Li (1997) solve additional cases by numerical methods.

across bidders, distributions have different shapes but approximately the same support.”

Maskin and Riley (2000b) also show quite generally that “strong” buyers prefer the second-price auction, whereas “weak” buyers prefer the first-price auction. This may be important where attracting buyers to enter the auction is an important consideration; see below.

1.7.2 Almost-Common-Values

If valuations have common-value components the effects of asymmetries can be even more dramatic. If one player has a small advantage, for example, a slightly higher private value in a setting that is close to pure-common-values, that player will bid a little more aggressively. This strengthens the opponent’s “winner’s curse” (since winning against a more aggressive competitor is worse news about the actual value of a common value object), so the opponent will bid a little less aggressively in an ascending auction, so the first player’s winner’s curse is reduced and she can bid a little more aggressively still, and so on. Klemperer (1998) discusses a range of contexts in which, in consequence, an apparently small edge for one player translates into a very large competitive advantage in an ascending auction. The earliest specific example in the literature is Bikhchandani’s (1988) demonstration that a small reputation advantage can allow a bidder to almost always win a pure-common-value auction, and that this reputational advantage may be very easy to sustain in a repeated context. Bulow, Huang, and Klemperer (1999) demonstrate that having a small toehold can be an enormous advantage in an otherwise pure-common-values takeover battle.⁵²

The original stimulus for all this work is Milgrom (1981)⁵³ which analyzes equilibria in ascending auctions and shows that there is a vast multiplicity in the pure-common values case, ranging from the symmetric equilibrium to equilibria in which an arbitrarily chosen player always wins. Later work shows that adding some “grit” into the model, whether it be a small private-value component, a small reputation component, or a small ownership component, etc., selects one of these equilibria, but which equilibrium is selected depends on exactly how the pure-common-values model is perturbed.

⁵² This analysis has been influential in competition policy. The UK Government recently blocked BSkyB (Rupert Murdoch’s satellite television company) from acquiring Manchester United (England’s most successful football club). An important reason was concern that by acquiring Manchester United, which receives the biggest share of the Premier League’s television revenues (about 7 percent), BSkyB would be able to shut out other television companies when the contract for the league’s broadcasting rights next comes up for auction (see *Economist*, March 20, 1999, p. 35; *Financial Times*, April 10, 1999, p. 22; and U.K. Monopolies and Mergers Commission, 1999).

⁵³ Also discussed in sections 1.7.3 and 1.8.3, and reprinted in *The Economic Theory of Auctions*.

Thus an apparently small change in the environment can greatly increase a player's chance of winning.

Since the winner of an "almost-common-value" ascending auction may therefore often have the lower signal, and so typically the lower marginal revenue, ascending auctions may be very unprofitable in this context.

By contrast, in a first-price auction a small change to the symmetric model generally results in a small change to the (unique) symmetric equilibrium, so the bidder with the higher signal hence, typically, higher marginal revenue continues to (almost always) win. Thus the first-price auction is almost optimal for a revenue-maximizing auctioneer, and is much more profitable than an ascending auction, provided bidders with higher signals have higher marginal revenues.⁵⁴

The effects of almost-common-values in ascending auctions are most extreme where there are also entry or bidding costs (see section 1.8) in which case the disadvantaged bidder(s) may not enter at all, leaving the auctioneer to face a single bidder (see Klemperer, 1998).

1.7.3 Information Advantages

Another important form of asymmetry is that one player may have superior information. Here, again, Milgrom (1981)⁵⁵ is critically important, showing that in a pure-common-value setting a bidder with no private information makes no profits in equilibrium in a second-price auction. Milgrom and Weber (1982b) show the same result (and much more) in the first-price context. The latter paper builds in part on work published later in Engelbrecht-Wiggans, Milgrom, and Weber (1983).

⁵⁴ However, Bulow and Klemperer (2002) show that the assumption that bidders with higher signals have higher marginal revenues is not innocuous in the common-values context. In the private-values context the assumption is equivalent to the assumption of downward-sloping marginal revenue for a monopolist whose demand corresponds to the distribution of a bidder's signals; in common-value settings, bidders' values and hence marginal revenues depend on *others'* signals, and oligopolists' marginal revenues are not necessarily decreasing in *other firms'* outputs. In the pure-common-values (and almost-common-values) cases the assumption is related to the assumption of strategic substitutes (see Bulow and Klemperer, 2002; and also Bulow, Geanakoplos, and Klemperer, 1985a,b).

Bulow and Klemperer (2002) show that if in fact this assumption does not hold, then a number of standard presumptions are violated in the symmetric equilibrium of a pure-common-values ascending auction; for example, more bidders can lower expected profits (see note 63) selling more units can raise average price, and rationing (as in Initial Public Offerings) can raise expected price. Furthermore, even if the assumption on marginal revenue holds, these results arise in the almost-common values case.

⁵⁵ Also discussed in sections 1.7.2 and 1.8.3, and reprinted in *The Economic Theory of Auctions*.

1.8 ENTRY COSTS AND THE NUMBER OF BIDDERS

1.8.1 *Endogenous Entry of Bidders*

In practical auction design, persuading bidders to take the time and trouble to enter the contest is a major concern, so we now endogenize the number of bidders and ask how it depends on the selling mechanism chosen.⁵⁶ (See also the Afterword to this chapter and parts C and D.)

The first key result is that in a private-value setting that satisfies the revenue-equivalence assumptions except for the presence of entry costs, bidders make the socially correct decision about whether or not to enter any standard auction if the reserve price is set equal to the seller's valuation. To see this, note that the expected social value of the bidder being present is the probability she wins the auction times the difference between her value and the runner-up's value. But this is exactly the bidder's expected profit after entering a second-price auction, and so also, using revenue equivalence, in a very wide class of auctions including all standard auctions.

Furthermore, in a free-entry equilibrium in which ex-ante identical bidders enter to the point at which each expects zero profits net of the entry cost (and each finds out her private value subsequent to the entry decision⁵⁷), the seller obtains the entire social surplus in expectation. So it also follows that running any standard auction with a reserve price equal to the seller's cost is revenue maximizing for the seller.⁵⁸

These results can be found in, for example, Levin and Smith (1994)⁵⁹ in a model in which bidders simultaneously make symmetric mixed-strategy entry decisions so that their expected profits are exactly zero. The results apply whether or not bidders observe how many others have chosen to enter before bidding in the auction, since revenue equivalence applies across and between both cases.⁶⁰ The results also apply if entry is sequential but the number of

⁵⁶ The Praetorians, when auctioning the Empire (see note 21), seem to have stipulated that the winning bidder could not punish the losers. This provision may have encouraged entry to the auction, although it would presumably reduce revenue from any exogenously fixed number of bidders.

⁵⁷ See Menezes and Monteiro (2000) for the case in which bidders know their private values prior to the entry decision.

⁵⁸ But the seller can increase social surplus, and hence her own expected revenue, if she can run a series of auctions. For example, she might announce an auction with a reserve price and the proviso that if the reserve is not met there will be a subsequent auction. Then there will be additional entrants in the second round if the good is not sold in the first round, that is, in the states in which the first-round entrants turned out to have low valuations. This increases social efficiency and seller revenue. See Burguet and Sákovics (1996) and also McAfee and McMillan (1988).

⁵⁹ Earlier related literature includes Engelbrecht-Wiggans (1987) and McAfee and McMillan (1987c). Levin and Smith (1996b) consider the seller's preference between standard auction forms when buyers are risk-averse; the auctioneer typically, but not always, prefers first-price to second-price auctions when bidders have entry costs.

⁶⁰ See section 1.8.4 below.

firms is treated as a continuous variable so that entrants' expected profits are exactly zero. In other cases the fact that the number of entrants must be an integer means that the marginal entrant's expected profits may exceed zero, but Engelbrecht-Wiggans (1993) shows that this makes very little difference: the seller optimally adjusts the reserve price and/or sets an entry subsidy or fee that sucks up all the entrants' surplus while altering the number of entrants by at most one.^{61,62}

In pure-common-value auctions, in marked contrast, the socially optimal number of bidders is obviously just one. Furthermore, Matthews (1984) shows that expected seller revenue can also be *decreasing* in the number of bidders in non-pathological pure-common-value settings, and Bulow and Klemperer (2002) provide intuition for why this is natural by using marginal-revenue analysis.⁶³ So both socially and privately, entry fees and reservation prices are much more desirable in common-value contexts. See Levin and Smith (1994) and also Harstad (1990).

Where bidders are asymmetric ex-ante, an auctioneer may wish to run an ex-post inefficient auction to attract weaker bidders to enter the contest. Thus Gilbert and Klemperer (2000) show that committing to ration output (i.e., selling at a fixed price at which demand exceeds supply) may be more profitable than raising price to clear the market (i.e., running an ascending auction that is ex-post efficient) because it attracts more buyers into the market.^{64,65}

Finally, bidders can influence the number of their rivals through their own strategic behavior. In particular, Fishman (1988) demonstrates that it can be profitable for a bidder to commit to a high bid (e.g., by making a preemptive

⁶¹ The same result extends to affiliated private value auctions.

⁶² Furthermore, Bulow and Klemperer (1996) show in the same setting that an additional bidder is worth more than the ability to set an optimal reserve price against a given number of bidders. See section 1.8.2.

⁶³ The point is that while the assumption that bidders with higher signals have higher marginal revenues usually holds in private-value settings, it often does not hold in pure-common-value settings. See note 54.

For a simple example consider the case where the common value equals the maximum of three signals v_i , $i = 1, 2, 3$, each drawn independently from a uniform distribution on $[0, 1]$ and each known by a different bidder. By selling to an uninformed bidder the seller makes $\frac{9}{12}$ ($= E\max\{v_1, v_2, v_3\}$). Selling to a single informed bidder the maximum revenue equals $\frac{8}{12}$ ($= E\max\{v_2, v_3\}$) achieved by setting a reservation price at this level (an informed bidder with signal 0 will just be willing to pay this price). Selling to two informed bidders yields at most $\frac{7}{12}$ in expectation (by a slightly harder calculation). Selling to all three bidders yields at most $\frac{6}{12}$ in expectation (the expected second-highest signal). (See exercise 8.)

⁶⁴ Or because it lowers the cost of persuading a given number of buyers to invest in participating in the market. Possible examples include the rationing of microchips, and split-award defense contracts.

Bulow and Klemperer (2002) provide another reason for rationing. See the previous paragraph.

⁶⁵ Similarly, Persico (2000b) shows that bidders have more incentive to collect information prior to a first-price auction than prior to a second-price auction.

“jump” bid in a takeover battle) to deter potential rivals from incurring the cost required to enter the contest.⁶⁶

1.8.2 The Value of Additional Bidders

Bulow and Klemperer (1996) show that when bidders are symmetric, an additional bidder is worth more to the seller in an ascending auction than the ability to set a reserve price, provided bidders with higher signals have higher marginal revenues. They then demonstrate that, very generally in a private-value auction, and also in a wide class of common-value settings,⁶⁷ a simple ascending auction with no reserve price and $N + 1$ symmetric bidders is more profitable than any auction that can realistically be run with N of these bidders.⁶⁸ So it is typically worthwhile for a seller to devote more resources to expanding the market than to collecting the information and performing the calculations required to figure out the best mechanism.

1.8.3 Information Aggregation with Large Numbers of Bidders

An important strand of the auction literature has been concerned with the properties of pure-common-value auctions as the number of bidders becomes large. The question is: does the sale price converge to the true value, thus fully aggregating all of the economy’s information even though each bidder has only partial information? If it does, it is then attractive to think of an auction model as justifying some of our ideas about perfect competition.

Wilson (1977)’s important early contribution showed that the answer is “yes” for a first-price auction under his assumptions.⁶⁹ Milgrom (1981) obtained similar results for a second-price auction (or for a $(k + 1)$ th price auction for k objects) in his remarkable paper that contains a range of other significant results and which we have already mentioned in sections 1.7.2 and 1.7.3.⁷⁰

Matthews (1984) allows each bidder to acquire information at a cost. In his model, as the number of bidders becomes large the amount of information each obtains falls, but in such a way that the (first-price) sale price does not in general converge to the true value.

⁶⁶ Similarly, Daniel and Hirshleifer (1995) obtain jump bidding in an ascending auction when each successive bid is costly. See section 1.13.3.

⁶⁷ In common-value settings higher signals would not imply higher marginal revenues. See note 54.

⁶⁸ A Crémer and McLean (1985)-style mechanism is probably *not* a realistic one in a practical setting. See our discussion in section 1.6, and also Bulow and Klemperer (1996) and Lopomo (1998).

⁶⁹ Milgrom (1979) gives a precise characterization of the signal structures that imply Wilson’s result.

⁷⁰ Pesendorfer and Swinkels (1997) show that the sale price converges to the true value in a $(k + 1)$ th price auction for k objects under weaker assumptions than Wilson’s, provided that the number of objects as well as the number of bidders becomes large.

1.8.4 Unknown Number of Bidders

Matthews (1987)⁷¹ and McAfee and McMillan (1987b) consider auctions when bidders with private values are uncertain about how many rivals they are competing with,⁷² and analyze how bidders' and the seller's preferences between the number of bidders being known or unknown depend on the nature of bidders' risk aversion and on whether bidders' signals are affiliated, etc.⁷³

Finally, it is not hard to see that under the usual assumptions (risk-neutrality, independent signals, symmetry, etc.), the standard argument for revenue equivalence applies independent of whether the actual number of competitors is revealed to bidders before they bid.⁷⁴

1.9 COLLUSION

A crucial concern about auctions in practice is the ability of bidders to collude, but the theoretical work on this issue is rather limited. (See, however, the Afterword to this chapter and parts C and D.)

Robinson (1985) makes the simple but important point that a collusive agreement may be easier to sustain in a second-price auction than in a first-price auction. Assuming, for simplicity, no problems in coming to agreement among all the bidders, or in sharing the rewards between them, and abstracting from any concerns about detection, etc., the optimal agreement in a second-price auction is for the designated winner to bid infinitely high while all the other bidders bid zero, and no other bidder has any incentive to cheat on this agreement. But to do as well in a first-price auction the bidders must agree that the designated winner bid an arbitrarily small amount, while all the others bid zero, and all the others then have a substantial incentive to cheat on the agreement.⁷⁵

An important question is whether the cartel can find a mechanism that efficiently (and incentive-compatibly) designates the winner and divides the

⁷¹ Also discussed in section 1.5, and reprinted in the corresponding part of *The Economic Theory of Auctions*.

⁷² In Piccione and Tan (1996) the number of bidders is known, but the number of bidders with private information is uncertain. This paper also considers common-value settings.

⁷³ McAfee and McMillan (1987b) also consider optimal auctions for the case of risk-neutral bidders.

⁷⁴ See Harstad, Levin, and Kagel (1990) for explicit bidding functions for the standard auction forms when the assumptions for revenue equivalence apply. The revenue equivalence result is also a special case of results in Matthews (1987) and McAfee and McMillan (1987b).

⁷⁵ Milgrom (1987) develops a similar intuition to argue that repeated second-price auctions are more susceptible to collusion than repeated first-price auctions.

spoils by making appropriate side payments, when bidders have private information about their own values. McAfee and McMillan (1992)'s main result is that this is possible and can be implemented by a simple pre-auction if all the bidders in the auction are members of the cartel and they all have private values drawn from the same distribution. This result is very closely related to the demonstration in Cramton, Gibbons, and Klemperer (1987)⁷⁶ that a partnership (e.g., the gains from a cartel) can be efficiently divided up.

McAfee and McMillan also analyze cartels that contain only a subgroup of the industry participants, and "weak cartels" that cannot make side payments between members, and consider how a seller should respond to the existence of a cartel.⁷⁷

Although there are many fewer formal analyses of collusion than seem merited by the issue's practical importance, Hendricks and Porter (1989) is a very useful informal survey of the circumstances and mechanisms facilitating collusion. They focus especially on methods of detecting collusion.

1.10 MULTI-UNIT AUCTIONS

Most auction theory, and almost all of the work discussed this far, restricts attention to the sale of a single indivisible unit. The literature on the sale of multiple units is much less well developed, except for the case where bidders demand only a single unit each. It is, however, the most active field of current research in auction theory,⁷⁸ so this is probably the section of this survey that will become obsolete most quickly. (See the Afterword to this chapter.)

1.10.1 Optimal Auctions

Maskin and Riley (1989) extend Myerson's (1981) analysis of optimal auctions to the case in which buyers have downward-sloping demand-curves,

⁷⁶ Discussed in section 1.12.2, and reprinted in that part of *The Economic Theory of Auctions*.

⁷⁷ Graham and Marshall (1987) address similar issues and show how any subset of bidders can achieve efficient collusion if an outside agent is available to achieve ex-post budget balancing (see also Graham, Marshall, and Richard, 1990). Mailath and Zemsky (1991) show how to achieve efficient collusion in second-price auctions, even among a subset of bidders who are not ex-ante identical and without the need for an outside agent, but using a more complicated mechanism. Hendricks, Porter, and Tan (1999) derive a necessary and sufficient condition for an efficient, incentive-compatible cartel in a common-value setting.

⁷⁸ Much current work has been stimulated by the recent government auctions of radio spectrum licenses (for mobile telephony, etc.), and emphasizes the problem of selling heterogeneous goods with complementarities between them, with common-value components to bidders' valuations, and perhaps also externalities between bidders. For discussion of the spectrum sales see McAfee and McMillan (1994, 1996), Klemperer (1998) (discussed in section 1.7.2, and reprinted in the corresponding part of *The Economics of Auctions*) and especially, Milgrom (2004). Another large body of important work has been stimulated by treasury auctions. See Bikhchandani and Huang (1993) for a survey of treasury security markets.

independently drawn from a one-parameter distribution, for quantities of a homogeneous good.⁷⁹ They provide one of a number of expositions of revenue equivalence for the multi-unit case, when buyers each want no more than a single unit.

Palfrey (1983) analyzes a seller's (and buyers') preferences between bundling heterogeneous objects and selling them unbundled; he shows the seller's incentive to bundle diminishes as the number of bidders increases. Very little progress has been made since Palfrey's paper on the problem of determining the seller-optimal auction for selling heterogeneous objects, but this topic is the subject of active current research.⁸⁰

1.10.2 Simultaneous Auctions

Wilson (1979), in another of his papers that was many years ahead of its time, first analyzed *share auctions*—auctions in which each bidder offers a schedule specifying a price for each possible fraction of the item (e.g., a certain volume of Treasury notes). He showed that in a uniform-price auction (when all the shares are sold at the (same) price that equates the supply and demand of shares) there are Nash equilibria that look very collusive, in that they support prices that may be much lower than if the item were sold as an indivisible unit. The intuition is that bidders can implicitly agree to divide up the item at a low price by each bidding extremely aggressively for smaller quantities than her equilibrium share so deterring the others from bidding for more.

This intuition suggests (at least) two ways of “undoing” the equilibrium. One way is to run a discriminatory auction in which bidders pay the price they bid for each share; bidding aggressively for small quantities is then very costly, so bidders submit flatter demand curves which induce greater price competition at the margin. See Back and Zender (1993), who argue that discriminatory auctions are therefore likely to be far more profitable for a seller.⁸¹ Nevertheless, Anton and Yao (1992) show that implicit coordination is still possible in this kind of auction if bidders' values are non-linear in the volume purchased.⁸²

A second way of undoing the low-price uniform-price equilibrium is to include some randomness in demands (e.g., from non-competitive bidders)

⁷⁹ As for Myerson (1981), the analysis can be interpreted through marginal revenues, though it is not presented this way.

⁸⁰ See, for example, Armstrong (2000) and Avery and Hendershott (2000), and Rothkopf, Pekeč, and Harstad (1998).

⁸¹ The section of Wilson's paper on discriminatory auctions is a little misleading about the relationship with uniform-price auctions.

Maxwell (1983) is earlier work extending Wilson's paper.

⁸² Anton and Yao also use a private-value framework in contrast to Back and Zender's and Wilson's common-value setting. See also Bernheim and Whinston (1986) and Anton and Yao (1989) for related models without incomplete information about costs or values.

or in the seller's supply. Klemperer and Meyer (1989) take this tack and show that sufficient supply uncertainty can reduce the multiplicity of uniform-price equilibria to a single equilibrium that is highly competitive if bidders' values are linear in their volumes purchased.⁸³ They pose their model in an oligopoly setting, or equivalently a procurement auction, and allow non-linear (but publicly-known) costs; the model closely corresponds to the actual operation of electricity-supply auction markets.⁸⁴

Klemperer and Meyer's model allows downward-sloping demand (in the procurement context) hence the quantity is endogenous to the bids (even absent demand uncertainty). Hansen (1988) considers endogenous quantity in the winner-take-all context, and shows that not only does the auctioneer prefer a first-price to a second-price auction (in a context where revenue equivalence would hold if the quantity were fixed) but the first-price auction is also socially more efficient and may even be preferred by the bidders. The intuition is that in first- and second-price auctions the quantity traded depends on the prices bid by the winner and the runner-up, respectively. So the first-price auction is more productively efficient (the quantity traded reflects the winner's cost or value) and provides greater incentive for aggressive bidding (a more aggressive bid not only increases the probability of winning, but also increases the quantity traded contingent on winning).

1.10.3 Sequential Auctions

The analysis of auctions where units are sold sequentially is well developed for the important special case in which no buyer is interested in more than one unit. In this case, if the units are homogeneous, and under the other usual assumptions, revenue equivalence holds whether the units are sold sequentially or simultaneously (Weber, 1983; Maskin and Riley, 1989; Bulow and Klemperer, 1994).

Thus quite complex multi-unit auctions can be solved by using revenue equivalence to work out, at any point of the game, what players' strategies must be to yield them the same expected payoff as if all the remaining units were auctioned simultaneously in a simple ascending auction.

Bulow and Klemperer (1994) use the revenue equivalence theorem in this way to solve for the dynamic price-path of a model of a stock market or housing market; the model would be intractably hard to solve by the direct method of writing down the first-order conditions for equilibrium in a dynamic

⁸³ Back and Zender (1993) argue that realistic amounts of uncertainty may nevertheless leave a continuum of equilibria. See Nyborg (1997) for further discussion and other arguments against the low-price equilibrium. Other related recent work on simultaneous multi-unit auctions includes Daripa (1996a,b), Engelbrecht-Wiggans and Kahn (1998a,b), and Wang and Zender (2002).

⁸⁴ See, for example, the developments of Klemperer and Meyer's model in Bolle (1992), Green and Newbery (1992), and Green (1996).

program. The point of the paper is that rational, strategic, traders *should* be very sensitive to new information and so participate in rushes of trading activity (frenzies) that sometimes lead to crashes in the market price. However, it is the method rather than the specific application that deserves emphasis here.

A much simpler example is the sale of k units through k repetitions of a first-price auction, with only the winning bid announced at each stage, to bidders with independent private values. Here, revenue equivalence tells us that at each stage each bidder just bids the expected $(k + 1)$ th highest value, conditional on being a winner and on the information revealed so far, since this is what she would pay if all the remaining units were auctioned simultaneously in a standard ascending auction. It is easy to see that this is a martingale, that is, the price neither rises nor falls over time, on average.

A large contribution of Milgrom and Weber's (2000) seminal paper is to consider a wider class of sequential auctions (including first-price auctions, both with and without price-announcements, second-price auctions, and English auctions) under more general assumptions. They show that with affiliation and/or common-value elements the price path drifts upwards. The intuition for the effect of affiliation is essentially that of the Linkage Principle (see section 1.6).⁸⁵ This paper has not previously been published, but it is a highly influential paper that it is gratifying to be able to publish at last in *The Economic Theory of Auctions*. Since it is unpolished, and the reader must beware of possible errors, we also publish a new foreword by the authors that explains the difficulties.

Milgrom and Weber's paper left a puzzle: contrary to the discussion above, it is more common to observe a *downward* drift in prices at auctions (see especially Ashenfelter, 1989). This discrepancy has spawned a small literature attempting to explain the "Declining Price Anomaly" (or "Afternoon Effect").⁸⁶ An early example is McAfee and Vincent (1993) who pursue the intuitive notion that risk-aversion might drive up early prices by providing an incentive to buy early. Actually, McAfee and Vincent's results are inconclusive; bidders use mixed strategies when risk-aversion is of the most plausible kind, so prices need not necessarily decline. Nevertheless, theirs is an impor-

⁸⁵ Note, however, that Perry and Reny (1998) show that the Linkage Principle need not hold if individuals can win more than one unit. The reason is that if (as in Milgrom and Weber's model) bidders desire at most one unit the underbidder is always a loser with pessimistic information, but in a multi-unit auction the underbidder for the marginal unit may already have won inframarginal units and have optimistic information.

⁸⁶ In fact Milgrom and Weber (2000) suggest a resolution of the "anomaly" themselves in their discussion of the 1981 sale of leases on RCA satellite-based telecommunications transponders. For other possible resolutions and analyses based on models in which no buyer demands more than one unit, see Bernhardt and Scoones (1994), Engelbrecht-Wiggans (1994), von der Fehr (1994), Gale and Hausch (1994), Beggs and Graddy (1997), and Ginsburg (1998). For analyses when bidders have multi-unit demand see several of the papers cited in note 88.

tant analysis and also provides an interesting example in which bidding functions that are monotonic in value do not exist.

Weber (1983) surveys many of the issues that arise in multi-object auctions, focusing primarily on sequential auctions. Unlike the previously mentioned papers in this subsection, he discusses the complex problems that arise when bidders desire multiple units; Ortega Reichert (1968b)⁸⁷ had already addressed some of these.^{88,89}

1.10.4 Efficient Auctions

A main message of much of the current research on multi-unit auctions is that it is very hard to achieve efficient outcomes.⁹⁰ This is in contrast to the single-unit case, in which Maskin (1992) showed under a broad class of assumptions that an ascending auction is efficient if bidders' private signals are single-dimensional, even with asymmetries among bidders and common-value components to valuations.

A Vickrey auction is efficient in private-value multi-unit contexts,⁹¹ and Dasgupta and Maskin (2000) and Perry and Reny (1998) show how to generalize the Vickrey mechanism to achieve efficiency in a wide variety of multi-unit contexts if each bidder's signal is one-dimensional. But Jehiel and Moldovanu (2001) obtain impossibility results showing that efficiency is not usually possible when each bidder's information signal is multidimensional, as is natural when there are multiple heterogeneous goods.

Ausubel (1998) and Ausubel and Cramton (1998a) emphasize the inefficiencies of standard auctions even in the sale of homogeneous objects. In particular, an ascending multi-unit auction (where the sale price equals the first price at which the number of units demanded falls to the supply available)

⁸⁷ Discussed in Weber's paper and in our section 1.2, and reprinted in the corresponding part of *The Economic Theory of Auctions*.

⁸⁸ Other nice papers analyzing sequential auctions when bidders have multi-unit demand include Robert's (≈ 1995) very elegant, tractable example; Pitchik and Schotter's (1988), Pitchik's (1995), and Benoît and Krishna's (2001) analyses of budget-constrained bidders; Levin's (1996), Gale, Hausch, and Stegeman's (2000), and von der Fehr and Riis's (1999) models of procurement auctions where bidders have increasing or decreasing marginal costs of supply, or capacity constraints, and the related analyses of Black and de Meza (1992), and Gale and Stegeman (2001); Krishna's (1993) application to whether incumbents will outbid potential entrants for capacity; and Hausch's (1986) analysis of sequential versus simultaneous sales in a model with some similarities to Ortega Reichert's.

⁸⁹ McAfee and Vincent (1997) consider an auctioneer who cannot commit not to re-auction an object that fails to meet its reserve, so who might hold multiple auctions of a single unit.

⁹⁰ This is true even in the complete information case (see Bikhchandani, 1999).

⁹¹ Although even in this context the Vickrey auction would be problematic for practical policy because high-valuers are often required to pay less than low-valuers (which seems odd to policy makers), because of the odd opportunities for collusive behavior, because of budget constraints, etc.

gives a large bidder an incentive to reduce her demand early in order to pay less for those units she does win.⁹²

1.11 ROYALTIES, INCENTIVES CONTRACTS, AND PAYMENTS FOR QUALITY

It is usually assumed that bidders' payments can depend only on the bids. But if the winner's value can be observed ex-post, even imperfectly, the seller can do better by making the winner's payment depend on this observation. This removes some of the winner's information rent, and can be interpreted as an application of the Linkage Principle.

Riley (1988) makes this point in a general context. As a practical application, the quantity of oil extracted may be a noisy signal of an oilfield's profitability; Riley shows that the seller's expected revenue can then be increased either by setting per-unit royalties that the winner must pay in addition to the fee bid, or by having bidders bid on the royalty rate they are willing to pay rather than on fixed fees.^{93,94}

Similarly, Laffont and Tirole (1987) analyze a procurement auction in which the winner will subsequently invest in unobserved effort to reduce its cost. The auctioneer observes the final realized cost. Auctioning an incentive contract with a cost-sharing provision gets a better price for the auctioneer by reducing the difference between firms' valuations of winning, so reducing the winner's rent (just like a royalty), even though it weakens the incentives for effort to reduce costs. One of Laffont and Tirole's key results is a "separation property": the optimal contract, and hence the winner's final cost, is similar to that which would apply if there were only a single firm and so no bidding competition, while the auction awards the contract to the firm that announces the lowest expected cost.⁹⁵

Che (1993) uses a version of Laffont and Tirole's model to analyze a multi-dimensional auction in which firms bid on both quality and price in a procurement auction. The auctioneer uses a scoring rule to evaluate the bids. It is no surprise that a revenue equivalence result applies, for example, between "first-score" and

⁹² Of course, the same effect is present in other models, for example, Klemperer and Meyer (1989).

⁹³ In analyzing this application, he builds on work by Reece (1979).

⁹⁴ But note royalties can be very dangerous in some settings. Imagine a government awarding a monopoly license for a market with downward-sloping demand to the firm that will pay the highest royalty per unit sold. Then firms with identical, constant, marginal costs will bid the royalty up to the vertical intercept of demand less this marginal cost. Government revenue, firm profits, and consumer surplus will all be zero. Riley assumes constant per unit revenue from the oil, and decreasing marginal cost up to some output level about which the bidders have private information.

⁹⁵ The same result is obtained independently in similar models due to McAfee and McMillan (1987d) and Riordan and Sappington (1987). A precursor to this work is McAfee and McMillan (1986).

“second-score” auctions. Che also shows that it is optimal for the auctioneer to pre-commit to a scoring rule that under-rewards quality relative to her real (ex-post) preferences.⁹⁶ Note that although this is a very elegant model of multi-dimensional bidding, firms only differ according to a one-dimensional type.⁹⁷

1.12 DOUBLE AUCTIONS, ETC.

1.12.1 Double Auctions

Standard auction theory assumes a single seller controls the trading mechanism, while many buyers submit bids. In a double auction, buyers and sellers are treated symmetrically with buyers submitting bids and sellers submitting asks. The double-auction literature thus provides a link to the bargaining literature. We emphasize here models that are closely related to simple, static, standard (one-sided) auctions.⁹⁸

The seminal model is the k -double auction of Chatterjee and Samuelson (1983) in which a single buyer and single seller submit a bid b and ask s , respectively, and if the bid exceeds the ask a trade is consummated at the price $kb + (1 - k)s$, where $0 \leq k \leq 1$. Of course, both buyer and seller have incentive to misrepresent their true values, so trades that would be efficient are not necessarily made.

Wilson (1985) first studied the generalization to the multi-buyer/multi-seller case in which each agent can trade at most one indivisible unit and, given the bids and asks, the maximum number of feasible trades are made at a price a fraction k of the distance between the lowest and highest feasible market clearing prices. The key result is that a double auction is efficient, in the sense that with sufficiently many buyers and sellers there is no other trading rule for which, conditional on agents' values it is common knowledge that all agents would be better off in expectation.

Rustichini, Satterthwaite, and Williams (1994) pursue the question of the extent to which agents' equilibrium bids and asks misrepresent their actual values. The answer is that in large markets the maximum misrepresentation is small, and hence the extent of the inefficiency caused by strategic behavior is also small.⁹⁹

Some intuition is provided by McAfee (1992) who considers the following simple mechanism: if N trades are feasible, let the $(N - 1)$ highest value

⁹⁶ In terms of the formal model, “quality” plays the role of the bidder's expected cost in Laffont and Tirole. Hence this result.

⁹⁷ See also Branco (1997) on multidimensional auctions.

⁹⁸ Furthermore, all the papers discussed in this section are independent private-value models. The assumption of private values, especially, seems important.

⁹⁹ Satterthwaite and Williams (1989a) and Williams (1991) had earlier obtained similar results for the special case $k = 1$ (the “buyer's bid double auction”) which is much easier to handle because sellers all have a dominant strategy of no misrepresentation.

buyers buy at the N th highest buyer's value while the $(N - 1)$ lowest value sellers sell at the N th lowest seller's value. Now, just as in a second-price auction, all agents report their actual values as dominant strategies, and only the least valuable feasible trade is foregone, and the mechanism also makes money. The fact that this mechanism is obviously so efficient (and McAfee shows how a slightly more complicated scheme does even better) makes it less surprising that other double auction mechanisms are also very efficient.

1.12.2 Related Two-Sided Trading Mechanisms

Related important work includes Myerson and Satterthwaite's (1983) path-breaking general analysis of mechanism design for bilateral trading. They use techniques similar to those of Myerson (1981), and the reader is similarly urged to study the reinterpretation in terms of marginal revenues and marginal costs given in Bulow and Roberts (1989).¹⁰⁰ Myerson and Satterthwaite show that the symmetric version ($k = \frac{1}{2}$) of Chatterjee and Samuelson's two-player double auction is in fact an optimal mechanism, in that it maximizes the expected gains from trade, in the case that the agents' values are independently drawn from identical uniform distributions.¹⁰¹

This paper also demonstrates that ex-post efficiency cannot be achieved in bargaining between a seller who initially owns the asset and a buyer with no prior ownership, when there is private information about valuations. However, Cramton, Gibbons, and Klemperer (1987) show that ex-post efficiency can be guaranteed (i.e., is consistent with incentive compatibility and individual rationality) when the asset to be traded is jointly owned: the reason is that traders' incentives to misrepresent their values are reduced by their uncertainty about whether they will be buyers or sellers. Cramton, Gibbons, and Klemperer exhibit one bidding game that achieves efficiency; revenue equivalence means that other auction forms can achieve the same outcome.¹⁰² This paper explains why ex-post efficient collusion in an auction (i.e., efficiently dividing the joint spoils by designating a winner and making appropriate side payments) is possible (see section 1.9).

1.13. OTHER TOPICS

This section considers some other important topics, each of which is represented by a paper in *The Economic Theory of Auctions*.

¹⁰⁰ These two papers are both discussed in section 1.4 and reprinted in the corresponding part of *The Economic Theory of Auctions*.

¹⁰¹ But this result depends critically on the distributional assumptions, and also assumes agents play the linear equilibrium constructed by Chatterjee and Samuelson. There are also non-linear equilibria (see Leininger, Linhart, and Radner, 1989; Satterthwaite and Williams, 1989b).

¹⁰² This paper, too, can be understood along the same lines that Bulow and Roberts explain Myerson and Satterthwaite.

1.13.1 Budget Constraints

An important reason why revenue equivalence may fail in practice is that bidders may face budget constraints. To see why, consider the standard model in which revenue equivalence applies and bidders have independent private values v_i , but let bidder i have budget constraint b_i . Then in a second-price auction i bids exactly as if she had value $x_i = \min(b_i, v_i)$ but no budget constraint, so by the revenue equivalence theorem¹⁰³ the expected revenue equals that from a first-price auction in which bidders have values x_i and no budget constraints, or equivalently a first-price auction in which bidders have values x_i and budget constraints x_i . It is intuitive that this is less expected revenue than from a first-price auction in which bidders have values v_i ($\geq x_i$) and budget constraints b_i ($\geq x_i$). So first-price auctions are more profitable than second-price auctions. This and similar results are obtained in Che and Gale (1998).

It is also intuitive that auction forms that take payments from losers, such as lotteries and “all-pay” auctions, can be more profitable still in the presence of budget constraints.¹⁰⁴ Budget constraints are also very important in sequential multi-unit auctions, where they provide incentives to, for example, try to reduce opponents’ budgets in early sales in order to lower subsequent sale prices. This is the subject of an important paper in the literature on experimental auctions, Pitchik and Schotter (1988), and is also an area of active research.¹⁰⁵

1.13.2 Externalities between Bidders

Jehiel and Moldovanu (1996) make an important extension to the theory by incorporating the possibility that a potential buyer cares who buys the object for sale in the event that she does not. This might be the case, for example, when a patent is auctioned to oligopolistic competitors, or when selling nuclear weapons.¹⁰⁶ Jehiel and Moldovanu’s paper raises many of the issues, including demonstrating in the context of first-price auctions with complete information that there may be multiple equilibria, and hence that a potential bidder may do better to avoid an auction rather than show up and risk galvanizing an enemy to win. Jehiel, Moldovanu, and Stacchetti (1996) address the issue of constructing optimal mechanisms, and

¹⁰³ We assume the b_i are independently drawn from a strictly increasing atomless distribution, so that the x_i correspond to independent draws from a strictly increasing atomless distribution.

¹⁰⁴ See Che and Gale (1996).

¹⁰⁵ See Pitchik (1995), Benoît and Krishna (2001), and Harford (1998).

Budget constraints also affect the risk that a successful bidder may go bankrupt, or otherwise fail to honor the sale contract. See Board (1999), Hansen and Lott (1991), Spulber (1990), Waehrer (1995), and Zheng (2001).

¹⁰⁶ Or an empire. See note 21.

Caillaud and Jehiel (1998) show that externalities between bidders tend to make collusion harder.¹⁰⁷

1.13.3 *Jump Bidding*

An ascending auction is usually modeled as a continuous process in which each successive bid is an arbitrarily small increment above the previous bid. However, actual behavior in, for example, takeover battles, often involves “jump bidding” in which a bidder raises the price very substantially with a single bid.¹⁰⁸

To understand why this might happen, consider a standard independent private-value English auction with two symmetric players. The following behavior is an equilibrium: one player bids the price she would bid in a *first-price sealed-bid auction*; the second player then infers the first player’s actual value and bids that actual value if her own value is higher, but quits the auction otherwise.¹⁰⁹ So the player with the actual higher value wins, but the first player pays the first-price auction price when she wins, while the second player pays the second-highest valuation when she wins. Since the higher-value bidder always wins, the outcome is revenue-equivalent to that of the standard continuously ascending model in which the winner always pays the second-highest valuation. And since the first bidder may fear that the second may misunderstand the equilibrium and bid up the price when she will anyway lose, it is not the most natural equilibrium in the simple independent private-value model. But with affiliation, bidders prefer first-price sealed-bid auctions to continuous ascending auctions, as shown in Milgrom and Weber (1982a),¹¹⁰ so the first-price features of this equilibrium are attractive to bidders, and Avery (1998) demonstrates that we may therefore expect a “jumping” equilibrium to be played.

If there are costs to making each bid, then jump bidding arises for similar reasons, even with independent private values, see Daniel and Hirshleifer (1995).

1.13.4 *The War of Attrition*

The War of Attrition is no more than a special kind of auction in which all the bidders pay, and keep on paying at some specified rate, until they quit compet-

¹⁰⁷ In a related vein, Fullerton and McAfee (1999) examine bidders who are concerned about the risk of entering an industry against stronger rivals.

¹⁰⁸ The auction of the Empire (see note 21) was settled by a final jump bid from 5,000 drachms to 6,250 drachms, though in this case bidders’ strategies were probably not optimal ex-ante (and certainly not ex-post).

¹⁰⁹ To my knowledge Daniel and Hirshleifer were the first to note that this kind of jump bidding is an equilibrium of the basic model even absent affiliation or bidding costs.

¹¹⁰ Discussed in section 1.6 and reprinted in the corresponding part of *The Economic Theory of Auctions*.

ing for the prize. (It is irrelevant to the analysis that in most practical contexts the payments are social waste, rather than collected by an auctioneer.) Important early contributions were made by Riley (1980), Bliss and Nalebuff (1984), and Fudenberg and Tirole (1986), among others. Bulow and Klemperer (1999) extends the analysis to the many-player case, and makes the auction-theoretic underpinnings most explicit, including several appeals to revenue equivalence arguments.^{111,112}

1.13.5 Competing Auctioneers

McAfee (1993) examines a model in which many sellers compete for buyers. In equilibrium, in an infinitely large market, each seller holds an efficient auction including setting an efficient reserve price. Thus McAfee and ensuing papers endogenize the use of auctions, and so address the question of when we should expect auctions to arise.¹¹³

1.14. TESTING THE THEORY

This chapter is concerned with the theory of auctions, but its companion book, *The Economic Theory of Auctions*, concludes with recent surveys of the empirical evidence, Laffont (1997),¹¹⁴ and of the experimental evidence, Kagel (1995).

1.15. CONCLUSION

Auction theory has been among the most successful branches of economics in recent years. The theory has developed rapidly, and is increasingly being looked to for assistance in practical applications. Testing auction-theoretic models is seen as one of the brightest spots in applied economics. Much

¹¹¹ Krishna and Morgan (1997) analyze an open-loop War of Attrition (i.e., bidders cannot revise their strategies based on others' drop-out times) and also analyze the closely related All-Pay Auction.

¹¹² See section 2.2.2. Many other models of tournaments, lobbying, political contests, R&D races, etc., can most easily be understood as auctions (see chapter 2).

¹¹³ See Peters and Severinov (1997) and Burguet and Sákovic (1999) for further developments along McAfee's lines. In related veins, Manelli and Vincent (1995) study when a procurement auction is more desirable than sequential negotiation, if potential suppliers are privately informed about their goods' qualities; and Bulow and Klemperer (1996) show a standard auction with a reserve price at the auctioneer's value is more profitable than any negotiating mechanism (or optimal auction) if the standard auction attracts at least one more participant (see section 1.8.2).

¹¹⁴ A good alternative is Hendricks and Paarsch (1995), while Porter (1995) and Laffont and Vuong (1996) offer valuable surveys covering a more limited range. Among the outstanding research articles are Hendricks and Porter (1988) and Laffont, Ossard, and Vuong (1995).

research remains to be done, especially perhaps on multi-unit auctions, and much research is currently being done. But the foundations of the subject, as presented in many of the papers described here, seem secure.

APPENDIX 1.A: THE REVENUE EQUIVALENCE THEOREM

For simplicity, we focus on the basic “independent private values” model, in which n bidders compete for a single unit. Bidder i values the unit at v_i , which is private information to her, but it is common knowledge that each v_i is independently drawn from the same continuous distribution $F(v)$ on $[\underline{v}, \bar{v}]$ (so $F(\underline{v}) = 0$, $F(\bar{v}) = 1$) with density $f(v)$. All bidders are risk-neutral.

Consider *any* mechanism (any single-stage or multi-stage game) for allocating the unit among the n bidders. For this mechanism, and for a given bidder i , let $S_i(v)$ be the expected surplus that bidder i will obtain in equilibrium from participating in the mechanism, as a function of her type, which we now denote by v , rather than v_i , for notational convenience. Let $P_i(v)$ be her probability of receiving the object in the equilibrium. So $S_i(v) = vP_i(v) - E(\text{payment by type } v \text{ of player } i)$.

The following equation is the key:

$$S_i(v) \geq S_i(\tilde{v}) + (v - \tilde{v})P_i(\tilde{v}). \quad (1)$$

The right-hand side is the surplus that player i would obtain if she had type v but deviated from equilibrium behavior, and instead followed the strategy that type \tilde{v} of player i is supposed to follow in the equilibrium of the game induced by the mechanism. That is, if type v exactly mimics what type \tilde{v} would do, then v makes the same payments and wins the object as often as \tilde{v} would. So v gets the same utility that \tilde{v} would get ($S_i(\tilde{v})$), except that in states in which \tilde{v} would win the object (which happens with probability $P_i(\tilde{v})$) type v values the object at $(v - \tilde{v})$ more than \tilde{v} does, and so v obtains an extra $(v - \tilde{v})P_i(\tilde{v})$ more surplus in all. In an equilibrium, v must prefer not to deviate from equilibrium behavior, so the left-hand side must (weakly) exceed the right-hand side.

So, since type v must not want to mimic type $v + dv$, we have

$$S_i(v) \geq S_i(v + dv) + (-dv)P_i(v + dv) \quad (2)$$

(this is just (1) with \tilde{v} substituted by $v + dv$), and since $v + dv$ must not want to mimic type v we have

$$S_i(v + dv) \geq S_i(v) + (dv)P_i(v). \quad (3)$$

Reorganizing (2) and (3) yields

$$P_i(v + dv) \geq \frac{S_i(v + dv) - S_i(v)}{dv} \geq P_i(v)$$

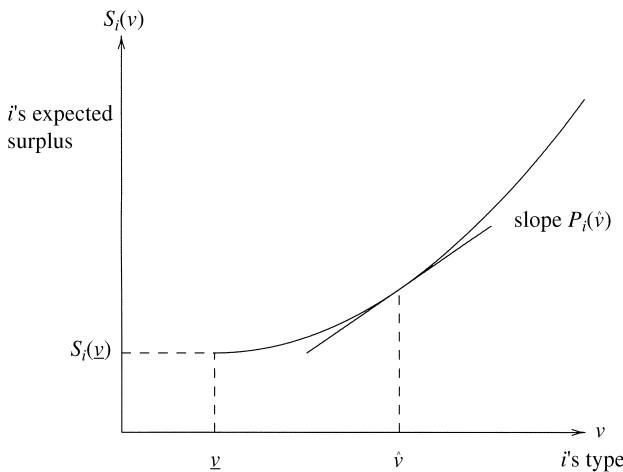


Figure 1.1 Bidder i 's expected surplus as a function of her type.

and taking the limit as $dv \rightarrow 0$ we obtain¹¹⁵

$$\frac{dS_i}{dv} = P_i(v). \tag{4}$$

Integrating up,

$$S_i(v) = S_i(\underline{v}) + \int_{x=\underline{v}}^v P_i(x)dx. \tag{5}$$

Equation (5) gives us a picture like Figure 1.1.

At any type \hat{v} the slope of the surplus function is $P_i(\hat{v})$, so if we know where the surplus function starts (i.e., know $S_i(\underline{v})$) we know the entire picture.

Now consider any two mechanisms which have the same $S_i(\underline{v})$ and the same $P_i(v)$ functions for all v and for every player i . They have the same $S_i(v)$ functions. So any given type, v , of player i makes the same expected payment in each of the two mechanisms (since $S_i(v) = vP_i(v) - E(\text{payment by type } v \text{ of player } i)$, since the bidder is risk-neutral). This means i 's expected payment averaged

¹¹⁵ An alternative way of obtaining this equation is to write $S_i(v) = T(v, \tilde{v}(v))$ equals v 's surplus when she behaves optimally as type $\tilde{v}(v)$. (In fact, $\tilde{v}(v) = v$.) Then the envelope theorem implies

$$\frac{dS_i}{dv} = \frac{dT}{dv} = \frac{\partial T}{\partial v}.$$

(That is,

$$\frac{dT}{dv} = \frac{\partial T}{\partial v} + \frac{\partial T}{\partial \tilde{v}} \frac{d\tilde{v}}{dv}$$

but $\partial T / \partial \tilde{v} = 0$ when \tilde{v} is chosen optimally.)

But $\partial T / \partial v = P_i(v)$ since if the bidder's behavior is unchanged, the incremental utility from a value dv higher is $P_i(v)dv$.

across her different possible types, v , is also the same for both mechanisms. Since this is true for all bidders, i , the mechanisms yield the same expected revenue for the auctioneer.^{116,117}

This is the Revenue Equivalence Theorem. There are many different statements of it, but they all essentially give the results of the preceding paragraph in a more or less special form.

In particular any mechanism which always gives the object to the highest-value bidder in equilibrium (all the standard auction forms do this) has $P_i(v) = (F(v))^{n-1}$ (since a bidder's probability of winning is just the probability that all the other $(n - 1)$ bidders have lower values than she does), and many mechanisms (including all the standard ones) give a bidder of the lowest feasible type no chance of any surplus, that is, $S_i(\underline{v}) = 0$, so all these mechanisms will yield the same expected payment by each bidder and the same expected revenue for the auctioneer.

Notice that nothing about this argument (except the actual value of $P_i(v)$) relied on there being only a single object. Thus the theorem extends immediately to the case of $k > 1$ indivisible objects being sold, provided bidders want no more than one object each; all mechanisms that give the objects to the k highest-value bidders are revenue-equivalent. So we have:

Revenue Equivalence Theorem (Private-Value Case). *Assume each of n risk-neutral potential buyers has a privately known value independently drawn from a common distribution $F(v)$ that is strictly increasing and atomless on $[\underline{v}, \bar{v}]$. Suppose that no buyer wants more than one of the k available identical indivisible objects. Then any auction mechanism in which (i) the objects always go to the k buyers with the highest values, and (ii) any bidder with value \underline{v} expects zero surplus, yields the same expected revenue, and results in a buyer with value v making the same expected payment.*

¹¹⁶ Some readers may wish to think of this analysis in terms of the Revelation Principle (see Myerson, 1979; Dasgupta, Hammond, and Maskin, 1979; Harris and Townsend, 1981) that says that we can always restrict attention to direct revelation mechanisms that satisfy incentive compatibility. That is, any mechanism is equivalent to another mechanism in which agents report their types, v , and wish to do so truthfully. Here we have analyzed *any* auction by focusing attention on the equivalent truthful direct revelation mechanism. In our problem the incentive compatibility (truth-telling) constraints, (1), completely pin down the expected payments that must be made to each type of agent once $P_i(v)$ and $S_i(v)$ have been specified.

¹¹⁷ Note that this argument assumes that the distribution of types of bidder, v , has positive density everywhere on $[\underline{v}, \bar{v}]$ so that $dS_i(v)/dv$ is defined everywhere on the range, and hence $S_i(v)$ is completely determined by $S_i(\underline{v})$ and $P_i(v)$.

For example, assume instead that there are just two types, $v = 0$ and $v = 1$, and each of two bidders is equally likely to be of either type (independent of the other's type) and the seller begins by simultaneously offering both bidders a price α ; if just one accepts then the trade is made at price α , if both accept then the unit is allocated by lottery at price α , if neither accepts then the unit is allocated by lottery at price 0. Then a "high" type prefers to accept so $S_i(0) = 0$, $P_i(0) = \frac{1}{4}$, and $P_i(1) = \frac{3}{4}$, for any $\alpha \in (0, \frac{2}{3})$, but the seller's expected revenue is strictly increasing in α , so revenue equivalence fails. See also Harris and Raviv (1981), or closely study Maskin and Riley (1985). (See exercise 2.)

It is not hard to extend the result to the general (common- and/or private-value) case, in which each buyer, i , independently receives a *signal* t_i drawn from $[\underline{t}, \bar{t}]$ and each bidder's value $V_i(t_1, \dots, t_n)$ depends on all the signals.¹¹⁸ A more general statement of the theorem is then exactly the statement above, but with “signal” substituted for “value”, and t , \underline{t} , and \bar{t} substituted for v , \underline{v} , and \bar{v} , throughout.

Application to Computing Bidding Strategies

Again we focus, for simplicity, on the single-object private-value case.

One of the mechanisms satisfying the revenue equivalence theorem is the ascending auction, in which the expected payment of a bidder of type v is just $P_i(v)$ times the expectation of the highest of the remaining $(n - 1)$ values conditional on all these values being below v . Since the density of the highest of $(n - 1)$ values is $(n - 1)f(v)(F(v))^{n-2}$, this last expectation can be written as

$$\frac{\int_{x=\underline{v}}^v x(n - 1)f(x)(F(x))^{n-2} dx}{\int_{x=\underline{v}}^v (n - 1)f(x)(F(x))^{n-2} dx}$$

which, after integrating the numerator by parts,¹¹⁹ yields

$$v - \frac{\int_{x=\underline{v}}^v (F(x))^{n-1} dx}{(F(v))^{n-1}}.$$

Since in a first-price sealed-bid auction, v 's expected payments are $P_i(v)$ times her bid, it follows that v bids according to

$$b(v) = v - \frac{\int_{x=\underline{v}}^v (F(x))^{n-1} dx}{(F(v))^{n-1}}$$

in a first-price auction.

In an “all-pay” auction in which every competitor always pays her bid (but only the highest-payer wins the object), it likewise follows that v must bid

$$b(v) = (F(v))^{n-1} v - \int_{x=\underline{v}}^v (F(x))^{n-1} dx.$$

¹¹⁸ See, for example, Lemma 3 of Bulow and Klemperer (1996) (reprinted in *The Economic Theory of Auctions*).

¹¹⁹ The denominator integrates to $(F(v))^{n-1}$, and the numerator $\int_{x=\underline{v}}^v x[(n - 1)f(x)(F(x))^{n-2}] dx = \int_{x=\underline{v}}^v x d(F(x))^{n-1} = [x \cdot (F(x))^{n-1}]_{\underline{v}}^v - \int_{x=\underline{v}}^v 1 \cdot (F(x))^{n-1} dx$

which yields the result.

Computing the bidding strategies this way is somewhat easier than solving for them directly in these cases.^{120,121}

In other cases, see, for example, Bulow and Klemperer (1994),¹²² it is very much easier.¹²³

APPENDIX 1.B: MARGINAL REVENUES

This appendix develops the basics of the “marginal revenue” approach to auctions.

We begin by following Bulow and Klemperer (1996)¹²⁴ to show, very generally, that the expected revenue from an ascending auction equals the expected marginal revenue of the winning bidder.

Figure 1.2 plots value, v , against $1 - F(v)$ for bidder i . We can interpret this as a “demand curve” because bidder i ’s value exceeds any v with probability $1 - F(v)$, so if a monopolist faced the single bidder, i , and set a take-it-or-leave-it offer of price \hat{v} , the monopolist would make a sale with probability $1 - F(\hat{v})$, that is, the monopolist’s expected quantity of sales would be $q(\hat{v}) = 1 - F(\hat{v})$.

Figure 1.2 also shows a “marginal revenue curve”, $MR(v)$, constructed from

¹²⁰ Appendix 1.D illustrates for the case in which $F(\cdot)$ is uniform.

¹²¹ To solve directly for the first-price equilibrium bidding strategies, we look for a symmetric Nash equilibrium in which a bidder with value v chooses the bid $b(v)$, and assume (as can be proved, see, e.g., Example 6.5 of Fudenberg and Tirole (1991)) that b is a continuous strictly increasing function of v . Imagine player i with value v deviates and chooses the bid \tilde{b} . Let \tilde{v} be the type of bidder she would just tie with, that is, let $b(\tilde{v}) = \tilde{b}$. Mimicking \tilde{v} would beat all the other $(n - 1)$ bidders with probability $(F(\tilde{v}))^{n-1}$ and so yield expected surplus to player i of $T(v, \tilde{v}) = (v - b(\tilde{v}))(F(\tilde{v}))^{n-1}$.

Choosing the best bid to make is equivalent to choosing the best \tilde{v} to mimic, which we can do by looking at the first-order condition

$$\frac{\partial T(v, \tilde{v})}{\partial \tilde{v}} = -b'(\tilde{v})(F(\tilde{v}))^{n-1} + (v - b(\tilde{v}))(n - 1)(F(\tilde{v}))^{n-2}f(\tilde{v}).$$

For the bidding function $b(v)$ to be an equilibrium, i ’s best response to all others bidding according to this function must be to do likewise, that is, her optimal choice of \tilde{b} is $b(v)$ and of \tilde{v} is v . So

$$\frac{\partial T}{\partial \tilde{v}}(v, \tilde{v}) = 0 \quad \text{at } \tilde{v} = v \quad \Rightarrow \quad b'(v) = (v - b(v))(n - 1)\frac{f(v)}{F(v)}.$$

This differential equation can be solved for the equilibrium, using the boundary condition $b(\underline{v}) = \underline{v}$ (it is obvious type \underline{v} will not bid more than \underline{v} , and we assume the auctioneer will not accept lower bids than \underline{v}).

¹²² Described in section 1.10.3 and reprinted in the corresponding part of *The Economic Theory of Auctions*.

¹²³ For examples of using the revenue equivalence theorem to solve an oligopoly pricing problem, see Appendix B of Bulow and Klemperer (1998) and section 2.5.

¹²⁴ Also discussed in section 1.8.2, and reprinted in the corresponding part of *The Economic Theory of Auctions*.




Figure 1.2 “Demand” and “marginal revenue” curves for bidder with value v drawn from $F(v)$.

the demand curve in exactly the usual way.¹²⁵ Note that at “price” \hat{v} the monopolist’s expected revenues can be computed either as the horizontally shaded rectangle $\hat{v} \cdot [1 - F(\hat{v})]$, or as the vertically shaded area under the “marginal revenue” curve $MR(v)$ up to “quantity” $1 - F(\hat{v})$. That is, just as in standard monopoly theory, the monopolist’s revenues can be computed either as price times quantity, or as the sum of the marginal revenues of all the units sold. (Mathematically, we can write $\hat{v} \cdot [q(\hat{v})] = \int_{q=0}^{q(\hat{v})} MR(v(q))dq$.¹²⁶)

Now imagine bidder i is the winner of the ascending auction. Let \hat{v} be the actual value of the second-highest bidder. So the actual price in the auction equals \hat{v} . But the result of the previous paragraph tells us that \hat{v} equals the

¹²⁵ That is, “revenue” = “price” times “quantity” = $v \cdot q(v)$, so

$$MR(v) = \frac{d}{dq} [v \cdot q(v)] = v + \frac{q}{dq/dv} = v - \frac{1 - F(v)}{f(v)}.$$

¹²⁶ We can confirm

$$\int_{q=0}^{q(\hat{v})} MR(v(q))dq = \int_{v=\hat{v}}^{\bar{v}} MR(v)f(v)dv$$

(changing variables $q \rightarrow [1 - F(v)]$, $dq \rightarrow [-f(v)dv]$, reversing limits to cancel the minus sign and defining $q(\bar{v}) = 0$ as in Appendix 1.A) which equals

$$\int_{v=\hat{v}}^{\bar{v}} (vf(v) - 1 + F(v))dv = [vF(v) - v]_{v=\hat{v}}^{\bar{v}} = \hat{v} - \hat{v}F(\hat{v})$$

as claimed.

average level of the marginal revenue curve between 0 and $q(\hat{v})$. Mathematically, we have

$$\hat{v} = \frac{1}{q(\hat{v})} \int_{q=0}^{q(\hat{v})} MR(v(q))dq.$$

That is, \hat{v} equals the average value of i 's marginal revenue, conditional on i 's value exceeding \hat{v} . But what we know about i 's value is just that it exceeds \hat{v} , because i won the auction. So for any actual second-highest value \hat{v} , the price, and hence the actual revenue, equals the expected marginal revenue of the winner. *So the expected revenue from an ascending auction equals the expected marginal revenue of the winning bidder.*

Observe that the result is very general for ascending auctions. Nothing in the argument relies on bidders' private values being independent, nor on bidders being risk-neutral, nor on their values being drawn from a common distribution. It is also not hard to check that the argument extends directly to the general (common- and/or private-value) case. (See Bulow and Klemperer (1996) for full details.)

Obviously the result also extends to any auction that is revenue equivalent to the ascending auction. Noting the conditions for revenue equivalence (see Appendix 1.A) it follows that¹²⁷ *if the bidders are risk-neutral and their information signals are independent, the expected revenue from any standard auction equals the expected marginal revenue of the winning bidder.*

Alternative, Algebraic Proof

For the risk-neutral, independent, symmetric, private-value case we can alternatively obtain the result using the results (and notation) of Appendix 1.A:¹²⁸ Bidder i 's expected payment to the auctioneer equals i 's expected gross value received from the auction, $vP_i(v)$, less her expected surplus, $S_i(v)$. So the auctioneer's expected receipts from i are

$$\int_{v=\underline{v}}^{\bar{v}} (vP_i(v) - S_i(v))f(v)dv.$$

Substituting for $S_i(v)$ using equation (5) from Appendix 1.A yields

$$\int_{v=\underline{v}}^{\bar{v}} vP_i(v)f(v)dv - \int_{v=\underline{v}}^{\bar{v}} f(v) \int_{x=\underline{v}}^v P_i(x)dx dv - \int_{v=\underline{v}}^{\bar{v}} S_i(\underline{v})f(v)dv,$$

¹²⁷ It is not hard to check that bidders can be asymmetric, that is, their signals can be drawn from different distributions.

¹²⁸ This approach is the one taken by Bulow and Roberts (1989), who themselves follow Myerson (1981).

and integrating the second term by parts¹²⁹ yields

$$\int_{v=\underline{v}}^{\bar{v}} P_i(v)f(v)\left[v - \frac{1 - F(v)}{f(v)}\right]dv - S_i(\underline{v}).$$

Define bidder i 's "marginal revenue" if she has value v to be

$$MR_i(v) = \left[v - \frac{1 - F(v)}{f(v)}\right];$$

following the discussion of figure 1.2, above, this corresponds to thinking of v as "price", p , and of $1 - F(v)$ as "quantity", q , hence

$$\text{marginal revenue} = \frac{d(pq)}{dq} = p + \frac{q}{dq/dp} = v + \frac{1 - F(v)}{-f(v)}.$$

Then, assuming $S_i(\underline{v}) = 0$, as is the case for any standard mechanism (see Appendix 1.A), the auctioneer's receipts from all n bidders are

$$\sum_{i=1}^n \int_{v=\underline{v}}^{\bar{v}} P_i(v)f(v)MR_i(v)dv = \sum_{i=1}^n E_{v_i}[P_i(v_i)MR_i(v_i)]$$

in which, for convenience, we changed the dummy variable from v to v_i in the last expression. This expression equals the expected marginal revenue of the winning bidder. To see this, it is helpful to write $\tilde{P}_i(v_1, \dots, v_n)$ as the probability that i wins as a function of *all* bidders' signals (i.e., $\tilde{P}_i(v_1, \dots, v_n) = 1$ if i is the winner, $\tilde{P}_i(v_1, \dots, v_n) = 0$ otherwise). Then

$$P_i(v_i) = E_{v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n}[\tilde{P}_i(v_1, \dots, v_n)]$$

so expected total receipts can be written

$$E_{v_1, \dots, v_n} \left[\sum_{i=1}^n MR_i(v_i)\tilde{P}_i(v_1, \dots, v_n) \right]$$

which equals the expected marginal revenue of the winning bidder, since the term in square brackets is exactly the marginal revenue of the winner. If there is no sale for some realization of v_1, \dots, v_n , then the term in square brackets equals zero, so it is as if there was a winner with marginal revenue equal to zero.

So an auction that always sells to the bidder with the highest marginal revenue, except makes no sale if no bidder's marginal revenue exceeds zero, will maximize expected revenues. But with symmetric bidders, any standard auction will sell to the highest-value bidder. So if higher values imply higher marginal revenues, then any standard auction together with reserve price p_r such that $MR(p_r) = 0$

¹²⁹ That is,

$$\int_{v=\underline{v}}^{\bar{v}} f(v) \int_{x=\underline{v}}^v P_i(x)dx dv = \left[F(v) \int_{x=\underline{v}}^v P_i(x)dx \right]_{v=\underline{v}}^{\bar{v}} - \int_{v=\underline{v}}^{\bar{v}} F(v)P_i(v)dv = \int_{v=\underline{v}}^{\bar{v}} (1 - F(v))P_i(v)dv.$$

(to prevent any sale if all bidders have values below p_r , hence negative marginal revenues) maximizes the auctioneer's expected revenues.¹³⁰

Note also that this approach generalizes easily to bidders drawn from different distributions, and to the general (common- and/or private-value) case, but the risk-neutrality and independence assumptions are important, as they are for the revenue equivalence theorem.

APPENDIX 1.C: AFFILIATED SIGNALS

This appendix analyzes the relative profitabilities of the standard auction forms when bidders' signals are *affiliated*, illustrates how an optimal mechanism can extract the entire social surplus in this case, and provides a formal definition of affiliation.¹³¹

Loosely, two signals are affiliated if a higher value of one signal makes a higher value of the other signal more likely, *and* this is true on every subspace of the variables' domain. Thus affiliation is stronger than correlation which is just a global summary statistic; affiliation can be thought of as requiring local positive correlation everywhere.

Milgrom and Weber found that when signals are affiliated, the second-price open (i.e., ascending) auction raises more expected revenue than the second-price sealed-bid auction, which in turn beats the first-price auction (assuming risk-neutral bidders, whose signals are drawn from symmetric distributions, and whose value functions are symmetric functions of the signals). Why?

Begin with independent-private-value first-price auctions. Recall the intuition for equation (4) from Appendix 1.A,

$$\frac{dS_i(v)}{dv} = P_i(v),$$

that if a player has a value of $v + dv$ instead of v , she can emulate the strategy of a player with value v and win just as often, at the same cost, but earning an extra dv in the probability $P_i(v)$ event she wins.

Now consider affiliation. A player with a value of $v + dv$ who makes the same bid as a player with a value of v will pay the same price as a player with a value of v when she wins, but because of affiliation she will expect to win a bit less often. That is, her higher signal makes her think her competitors are also likely to have higher signals, which is bad for her expected profits.

But things are even worse in a second-price affiliated private-values auction for

¹³⁰ Appendix 1.D illustrates for the case of uniform $F(\cdot)$.

See Myerson (1981) and Bulow and Roberts (1989) for the design of revenue-maximizing auctions when bidders are asymmetric or when higher values do not always imply higher marginal revenues.

¹³¹ The first part of this appendix is based on notes written by Jeremy Bulow.

Appendix 1.D provides examples that illustrate the results.

the buyer. Not only does her probability of winning diminish, as in the first-price auction, but her costs per victory are higher. This is because affiliation implies that contingent on her winning the auction, the higher her value the higher the expected second-highest value which is the price she has to pay. Because the person with the highest value will win in either type of auction they are both equally efficient, and therefore the higher consumer surplus in the first-price auction implies higher seller revenue in the second-price auction.¹³²

How about second-price sealed-bid auctions versus ascending auctions? Sticking to private values, these two auction types will still be identical: the highest-valued bidder will always pay the second value. Also, with only two bidders there is no difference between sealed and open bids even with a common-value element and affiliation. In the open auction the player drops out when the price reaches her value for the good conditional on the other bidder having the same signal as her;¹³³ in the sealed-bid version a player bids her value conditional on the other player having the same signal.¹³⁴

If there are more than two bidders in a setting with affiliation and common values, then the ascending auction beats the sealed-bid auction: Assume there are three potential bidders in a second-price sealed-bid auction, each of whom reveals her signal to a trusted fourth party. The fourth party then tells the two most optimistic bidders that they are among the top two, but does not reveal the third bidder's signal. The first two will bid exactly as they would have without the information that they are in the top two, since their bids are only relevant in this case anyway. How will each bidder determine her bid? The marginal case in which it would be worthwhile for a bidder to win the auction is the case where she is tied for having the most optimistic signal. The second-highest actual bidder, whose bid determines the price, will bid the expectation of the asset's value, assuming that she is tied for the most optimistic assessment, and assuming there is a third observation with the distribution to be expected if, in fact, the second bidder *is* tied for the most optimistic signal.¹³⁵

¹³² To fill out this argument a little more, assume that in equilibrium there is some value for which the expected surplus (and therefore expected purchase price) is the same for a buyer in either type of auction. This will be true for the lowest-possible buyer value, for example, since that type of buyer obviously never wins the auction. Then by the argument in the text, the derivative of surplus with respect to value will be greater in the first-price auction than in the second-price auction. So the surplus in the first-price auction begins to grow faster, at least for a while, in the second-price auction. And if the surpluses ever came together again, the first-price surplus would have to forge ahead again. So on average across all possible bidder values, buyers will get more surplus in first-price auctions and sellers must therefore make more money in second-price auctions.

¹³³ To see why, check that if the other player is bidding this way then this player would lose money if she were to find herself a winner at a higher price (assuming higher signals imply higher values) but quitting at a lower price forgoes an opportunity to make money if the other player quits first. See Appendix 1.D for an example.

¹³⁴ The argument that the games are strategically equivalent is similar to the one for first-price and descending auctions.

¹³⁵ See Appendix 1.D for more discussion and examples.

However, the seller knows that contingent on the second bidder observing any given signal and there existing a first bidder with a more optimistic observation, the true distribution of this unknown third signal is a more optimistic one than the second bidder will use. (For example, given that the top two bidders have values of 30 and 40, the expectation of the third signal is higher than the expectation that the second bidder will use in a sealed-bid auction, which conditions on the top two values being 30 and 30.) Thus, the seller will do better on average to allow the bidder to make her offer contingent on the observation of the third bidder, as in the open auction where the third bidder's observation can be inferred from the price at which she drops out.

So, with affiliation, common-value elements, and more than two symmetric, risk-neutral, bidders, the first-price auction earns less revenue on average than the second-price sealed-bid auction which earns less than the ascending auction. With private values and/or only two bidders, the first-price auction still earns the least money but the other two types are tied.

Appendix 1.D provides some simple examples that illustrate these results.

Finally, we use a simple example to illustrate how an optimal mechanism can extract the entire social surplus from risk-neutral bidders whose signals are not independent. Let each of two bidders $i = 1, 2$ receive a private signal t_i which is either "high" or "low". Conditional on a bidder's signal, the other bidder receives the same signal with probability $\frac{2}{3}$ and receives the other possible signal with probability $\frac{1}{3}$. Bidder i 's actual value is $v_i(t_1, t_2)$. Now consider the following selling mechanism: (i) ask each bidder to report her signal; call these reports \tilde{t}_1 and \tilde{t}_2 respectively; (ii) if $\tilde{t}_1 = \tilde{t}_2$, pay each bidder an amount V ; (iii) if $\tilde{t}_1 \neq \tilde{t}_2$, require each bidder to pay $2V$ to the seller; and (iv) give the object to the bidder i with the highest value $v_i(\tilde{t}_1, \tilde{t}_2)$ at price $v_i(\tilde{t}_1, \tilde{t}_2)$. If V is sufficiently large, it is a Nash equilibrium for each bidder to "tell the truth", that is, report $\tilde{t}_i = t_i$ at stage 1, since if the other behaves this way, parts (ii) and (iii) of the mechanism then yield $\frac{2}{3}(V) + \frac{1}{3}(-2V) = 0$ to a truth-teller but yield $\frac{1}{3}(V) + \frac{2}{3}(-2V) = -V$ to a deviator. That is, the seller has essentially forced each bidder to make a bet on the other's signal, and the bidders can avoid losing money on these bets only by using their private information. But once all their private information has thus been revealed, the seller can extract the entire social surplus in part (iv) of the mechanism.¹³⁶

Formal Definition of Affiliation

Formally, but still restricting for simplicity to the case of two bidders, signals t_1 and t_2 are affiliated if for all $t_1' > t_1''$ and $t_2' > t_2''$,

¹³⁶ This simple mechanism suffers from additional equilibria that are not truth-telling—for example, for large V it is also an equilibrium for both bidders to always report "low"—but more complex mechanisms can be designed in which honesty is the unique equilibrium (see Myerson, 1981).

$$f(t'_1, t'_2)f(t''_1, t''_2) \geq f(t'_1, t''_2)f(t''_1, t'_2) \tag{6}$$

in which t'_i and t''_i are different possible realizations of the signals t_i , $i = 1, 2$, and $f(t_1, t_2)$ is the joint density function of the signals. Since, by the definition of conditional probability, $f(t_1, t_2) = g(t_1 | t_2)h(t_2)$, in which $g(t_1 | t_2)$ and $h(t_2)$ are the conditional density of t_1 given t_2 , and the unconditional density of t_2 , respectively, it follows that (6) holds if and only if

$$\frac{g(t'_1 | t'_2)}{g(t''_1 | t'_2)} \geq \frac{g(t'_1 | t''_2)}{g(t''_1 | t''_2)} \tag{7}$$

which is also known as the Monotone Likelihood Ratio Property, that is, higher values of t_1 become relatively more likely as t_2 increases. An implication of (7) is that¹³⁷

$$G(t_1 | t'_2) \leq G(t_1 | t''_2)$$

in which $G(t_1 | t_2)$ is the conditional distribution of t_1 given t_2 . In words, the distribution of t_1 conditional on t'_2 first-order stochastically dominates the distribution of t_1 conditional on t''_2 , if $t'_2 > t''_2$. The implications of affiliation that are probably used most frequently in auction-theory proofs are that¹³⁸

$$\frac{\partial}{\partial t_2} \left(\frac{g(t_1 | t_2)}{1 - G(t_1 | t_2)} \right) \leq 0$$

(in words, the hazard rate of t_1 is everywhere decreasing in t_2), and

$$\frac{\partial}{\partial t_2} \left(\frac{g(t_1 | t_2)}{G(t_1 | t_2)} \right) \geq 0$$

(the hazard rate of the decumulative density of t_1 is increasing in t_2 , that is, the probability density of $t = t_1$, conditional on $t \leq t_1$, is increasing in t_2).

In the case of independent signals, affiliation holds weakly.

APPENDIX 1.D: EXAMPLES USING THE UNIFORM DISTRIBUTION

This appendix uses the uniform distribution to develop some simple examples of bidding in the standard auctions, and illustrates the material in the preceding appendices.

The uniform distribution ($F(v) = (v - \underline{v})/(\bar{v} - \underline{v})$, $f(v) = (1/(\bar{v} - \underline{v}))$) is often particularly easy to work with in auction theory. The following fact is

¹³⁷ To obtain this, integrate first over $t'_1 > t_1$, then multiply through by $g(t''_1 | t'_2)g(t''_1 | t''_2)$, then integrate over $t''_1 < t'_1$ to yield an expression which implies this one.

¹³⁸ To obtain the next expression integrate (7) over $t'_1 > t''_1$, and substitute t_1 for t''_1 . To obtain the second expression, multiply (7) by $g(t''_1 | t'_2)g(t''_1 | t''_2)$, then integrate over $t''_1 < t'_1$ and substitute t_1 for t'_1 .

very helpful: the expected k th highest value among n values independently drawn from the uniform distribution on $[\underline{v}, \bar{v}]$ is

$$\underline{v} + \left(\frac{n+1-k}{n+1} \right) (\bar{v} - \underline{v}).$$

Bidding with Independent Private Values, Revenue Equivalence, and Marginal Revenues

Begin with risk-neutral bidders $i = 1, \dots, n$ each of whom has a private value v_i independently drawn from $[\underline{v}, \bar{v}]$.

Then in a second-price (or ascending) auction, everyone bids (or bids up to) her true value, so the seller's expected revenue is the expected second-highest value of the n values which, using the fact given above, equals

$$\underline{v} + \left(\frac{n-1}{n+1} \right) (\bar{v} - \underline{v})$$

In a first-price auction, by revenue equivalence, i bids her expected payment conditional on winning an ascending auction. Conditional on v_i being the highest value, the other $n-1$ values are uniformly distributed on $[\underline{v}, v_i]$ so, using the fact about the uniform distribution, the expected value of the highest of these—which is what i would expect to pay conditional on winning—is

$$\underline{v} + \left(\frac{n-1}{n} \right) (v_i - \underline{v})$$

so this will be i 's bid.

Alternatively, we can derive i 's equilibrium bidding strategy using the direct approach, and thus confirm revenue equivalence.¹³⁹

Note that the proportion of buyers with valuations above any price v is linear in v , that is,

$$q(v) \equiv 1 - F(v) = \frac{\bar{v} - v}{\bar{v} - \underline{v}}$$

Therefore $v = \bar{v} - (\bar{v} - \underline{v})q$, so the uniform distribution corresponds to linear demand (since v plays the role of price). It follows that the marginal revenue curve is just linear and twice as steep as the demand curve, that is,

$$MR(q(v)) = \bar{v} - 2(\bar{v} - \underline{v})q$$

¹³⁹ The differential equation for i 's first-price bidding strategy, $b(v_i)$, obtained using the direct approach, is $b'(v_i) = (v_i - b(v_i))(n-1)[f(v_i)/F(v_i)]$ (see note 121).

For the uniform distribution, this yields $b'(v_i) = (v_i - b(v_i))(n-1)[1/(v_i - \underline{v})]$ which is solved by $b(v_i) = \underline{v} + [(n-1)/n](v_i - \underline{v})$.

Since the highest-value bidder will determine the price, the seller's expected revenue will be $E\{\underline{v} + [(n-1)/n](\max_{i=1, \dots, n} v_i - \underline{v})\}$, so using our result that $E\{\max_{i=1, \dots, n} v_i\} = \underline{v} + [n/(n+1)](\bar{v} - \underline{v})$ yields that the expected revenue is $\underline{v} + [(n-1)/(n+1)](\bar{v} - \underline{v})$, confirming revenue equivalence with the second-price auction forms.

which implies¹⁴⁰

$$MR(v) = \bar{v} - 2(\bar{v} - v) \left(\frac{\bar{v} - v}{\bar{v} - \underline{v}} \right) = 2v - \bar{v}.$$

Since

$$E \left\{ \max_{i=1, \dots, n} v_i \right\} = \underline{v} + \left(\frac{n}{n+1} \right) (\bar{v} - \underline{v}),$$

the expected marginal revenue of the highest bidder equals

$$\underline{v} + \left(\frac{n-1}{n+1} \right) (\bar{v} - \underline{v}),$$

which confirms our earlier result that this is the expected revenue from any standard auction (without a reserve price).

Furthermore, since the marginal revenue curve is downward sloping, an optimal (i.e., expected-revenue maximizing) auction is any standard auction together with a reserve price, $p_r = \frac{1}{2}\bar{v}$ (so that $MR(p_r) = 0$), below which no sale will be made.

Bidding with Common Values, and the Winner's Curse

Now let the bidders have signals t_i , and $v_i = \alpha t_i + \beta \sum_{j \neq i} t_j$. (So $\beta = 0$ is the private values case, and $\alpha = \beta$ is pure common values; we assume $\alpha \geq \beta$.) Let $t_{(j)}$ be the actual j th highest signal.

In the symmetric equilibrium of an ascending auction each player quits where she would just be indifferent about finding herself a winner. So the first quit is at price $(\alpha + (n-1)\beta)t_{(n)}$, since that would be the actual value to all if all bidders had this signal; the remaining bidders all observe this and the next quit is at $\beta t_{(n)} + (\alpha + (n-2)\beta)t_{(n-1)}$ since this would be the current quitter's value if all the other remaining bidders were to quit with her; the other bidders all observe this and infer the next lowest signal, etc. The final quit, and so actual sale price is at

$$\hat{p} = \beta \sum_{j=3}^n t_{(j)} + (\alpha + \beta)t_{(2)}.$$

To check this is the equilibrium, note, for example, that if the player with the second-highest signal, $t_{(2)}$, waited to quit and found herself a winner at a price $\hat{p} + (\alpha + \beta)\varepsilon$ she would then infer $t_{(1)} = t_{(2)} + \varepsilon$ (since the final opponent is bidding symmetrically to her equilibrium behavior) hence that the value of the object to her was just $\hat{p} + \beta\varepsilon$, so she had lost money. But when

¹⁴⁰ Obviously, we can also obtain this equation by using the definition,

$$MR(v) = v - \frac{1 - F(v)}{f(v)},$$

from Appendix 1.B.

the price reached $\hat{p} - (\alpha + \beta)\varepsilon$, she could infer that her final opponent's signal is at least $t_{(2)} - \varepsilon$ hence that the value of the object to her was at least $\hat{p} - \beta\varepsilon$, so quitting early would have given up the opportunity of making some money (in the states in which the final opponent would have quit close to this price).¹⁴¹

Note that when the player with the second-highest signal quits, she knows (assuming equilibrium behavior) that the remaining signal is (weakly) higher than hers. So she is sure the actual value of the object to her cannot be less than the price at which she is quitting, and that the expected value is higher. This illustrates the *winner's curse*. The point is that what is relevant to her is *not* the expected value of the object, but rather its expected value conditional on her

¹⁴¹ The principle for solving the case where bidders' value functions are asymmetric is similar, and clarifies the argument. Assume just two bidders, for simplicity, with signals t_i, t_j and values $v_i(t_i, t_j)$ and $v_j(t_i, t_j)$. Assume that in equilibrium t_i quits at the same time as an opponent of type $t_j = w_i(t_i)$, in which $w_i(\cdot)$ is a strictly increasing function. So

$$b_i(t_i) = b_j(w_i(t_i)), \quad (*)$$

and t_i will beat all opponents with types $t_j < w_i(t_i)$, and lose to all higher types.

Now if t_i deviated from her equilibrium strategy and waited a tiny bit longer to quit, she would win against all $t_j \leq w_i(t_i)$ at the same prices as before, and she would also win against a few additional types of j with signals of (slightly above) $w_i(t_i)$, at a price of (slightly above) $b_i(t_i)$. Her value of winning in these additional cases would be (slightly above) $v_i(t_i, w_i(t_i))$, so if $b_i(t_i)$ were (strictly) less than $v_i(t_i, w_i(t_i))$, then deviating to win against a few additional types would be profitable. So $b_i(t_i) \geq v_i(t_i, w_i(t_i))$.

Similarly, if t_i were to quit a tiny bit earlier than her equilibrium quitting price, it would make no difference except that she would lose against a few types with signals (slightly below) $w_i(t_i)$ at prices of (slightly below) $b_i(t_i)$, and type t_i would wish to do this unless $b_i(t_i) \leq v_i(t_i, w_i(t_i))$.

So

$$b_i(t_i) = v_i(t_i, w_i(t_i)). \quad (**)$$

That is, t_i bids up to the value at which she would make no money if she were to find herself the winner.

Similarly

$$b_j(t_j) = v_j(w_j(t_j), t_j).$$

Substituting the value $t_j = w_i(t_i)$ into this equation yields

$$b_j(w_i(t_i)) = v_j(w_j(w_i(t_i)), w_i(t_i)).$$

But $b_j(w_i(t_i)) = b_i(t_i)$ by (*). And by definition $w_j(w_i(t_i)) = t_i$ (i.e., if t_i quits at the same time as $w_i(t_i)$, then the type that quits at the same time as $w_i(t_i)$ —this type is $w_j(w_i(t_i))$ —is t_i). So

$$b_i(t_i) = v_j(t_i, w_i(t_i)).$$

Comparing with (**) we have

$$v_i(t_i, w_i(t_i)) = v_j(t_i, w_i(t_i)).$$

That is, players have the same values when they have types that quit at the same time.

So to find the bidding strategies we solve this last equation for the function $w_i(t_i)$, and then substitute this function back into (**) to yield i 's bidding function. (Note that this procedure does not necessarily yield an equilibrium, although it does so in natural two-bidder or symmetric examples, see Maskin, 1992.)

winning it.¹⁴² Only when she wins the object does she care about its value, so she quits exactly at its value conditional on her winning. Exactly the same effect—that winning the auction is bad news about opponents’ signals, so bids must be adjusted down to allow for the “winner’s curse”—arises in the other auction types.

Note that the ascending auction equilibrium does not depend on the bidders’ signals being independent or on their distributions (which can be different for different bidders), or on the bidders being risk-neutral. However, these properties do not extend to the other standard auctions. So henceforth assume the signals are independent and uniform on $[0, \bar{t}]$, and the bidders are risk-neutral.

In a *second-price sealed-bid* auction the logic is similar to that for the ascending auction. Bidder i with signal t_i is willing to pay anything up to her expected value conditional on her winning the object but being just tied with one other with the same signal. The difference is that the bidder does not see the other $n - 2$ opponents’ bids, so estimates their signals at $\frac{1}{2}t_i$ (since conditional on them being below t_i , they are uniformly distributed below t_i). So i bids $\beta(n - 2)\frac{1}{2}t_i + (\alpha + \beta)t_i = (\alpha + \frac{1}{2}n\beta)t_i$.¹⁴³

The simplest way to solve for *first-price* bidding strategies is to use revenue equivalence.¹⁴⁴ Conditional on winning the second-price auction, a bidder with signal t_i expects to pay $(\alpha + \frac{1}{2}n\beta)\hat{t}$ in which \hat{t} is the expected highest of $n - 1$ signals uniformly distributed on $[0, t_i]$, that is,

$$\hat{t} = \left(\frac{n - 1}{n}\right)t_i.$$

So i bids this expected payment, that is,

$$\left(\frac{n - 1}{n}\right)(\alpha + \frac{1}{2}n\beta)t_i.$$

Bidding with Affiliated Signals, and Revenue Rankings

A tractable example of affiliated information that illustrates the revenue-ranking results derived in Appendix 1.C (and is also useful for developing other examples¹⁴⁵) has risk-neutral bidders $i = 1, \dots, n$ each of whom receives a signal t_i that is independently drawn from a uniform distribution on $[\nu - \frac{1}{2}, \nu + \frac{1}{2}]$ where ν is the (pure) common value of a single object for sale. Assume a “diffuse prior” for ν , that is, all values of ν are equally likely. (More formally we can let ν be uniformly distributed on $[-M, +M]$ and take the limit as $M \rightarrow \infty$.) So a higher value of t_i makes a higher value of ν more likely, and hence higher values of the other signals more likely,

¹⁴² This statement assumes risk-neutrality, but the point we are making obviously does not.

¹⁴³ We can also confirm this is the equilibrium either by revenue equivalence with the ascending auction, or by a similar argument to that for the ascending auction—the only effect on i of i bidding a small amount $(\alpha + \beta)\epsilon$ more is if i moves from coming second to winning, etc.

¹⁴⁴ An alternative is the direct method, see note 121.

¹⁴⁵ Most examples of affiliated information are very hard to work with.

and it can be checked that this example satisfies the formal definition of affiliation.

Let the j th highest actual signal be $t_{(j)}$, and observe that conditional on *all* the signals t_1, \dots, t_n , the expected value of v equals $\frac{1}{2}(t_{(1)} + t_{(n)})$ (since any value of $v \in [t_{(1)} - \frac{1}{2}, t_{(n)} + \frac{1}{2}]$ is equally probable).

We now compute the symmetric equilibria of the standard auction types.

In an *ascending auction*, the first quit will be at price $t_{(n)}$ (since that is where the lowest-signal bidder would be indifferent about winning were everyone else to quit simultaneously with her), and every other bidder i will then infer $t_{(n)}$ and quit at $\frac{1}{2}(t_{(n)} + t_i)$ (since that is where each i would be just indifferent about finding herself the winner). The price paid by the winner will therefore be $\frac{1}{2}(t_{(n)} + t_{(2)})$ which, using our result about the uniform distribution, on average equals

$$\frac{1}{2} \left(\left[v - \frac{1}{2} + \left(\frac{1}{n+1} \right) \right] + \left[v - \frac{1}{2} + \left(\frac{n-1}{n+1} \right) \right] \right) = v - \frac{1}{2} \left(\frac{1}{n+1} \right)$$

In a *sealed-bid second-price auction*, each bidder i bids her expected value, conditional on being tied for winner with one other bidder (see previous subsection). That is, i bids thinking of herself as being the highest of $n-1$ bidders uniformly drawn from $[v - \frac{1}{2}, v + \frac{1}{2}]$ and tied with one other, so on average, in this case,

$$t_i = \left[v - \frac{1}{2} + \left(\frac{n-1}{n} \right) \right]$$

so i 's estimate of v in this case, and hence her bid, equals

$$t_i + \frac{1}{2} - \left(\frac{n-1}{n} \right).$$

On average, the second-highest bidder of n bidders actually has signal

$$t_{(2)} = \left[v - \frac{1}{2} + \left(\frac{n-1}{n+1} \right) \right]$$

so bids

$$\left[v - \frac{1}{2} + \left(\frac{n-1}{n+1} \right) \right] + \left[\frac{1}{2} - \left(\frac{n-1}{n} \right) \right].$$

So the expected revenue from this auction equals

$$v - \left(\frac{n-1}{n} \right) \left(\frac{1}{n+1} \right).$$

In a *first-price auction*, likewise, each bidder i bids $t_i - x$ for some x ; this is because of our “diffuse prior” assumption which means that i 's signal gives her no information about whether she is high or low relative to others' signals or the

“truth”, and so should not affect how close she bids to her signal. Let $t_i = v - \frac{1}{2} + T_i$. In equilibrium i will have the highest signal, and so win the auction, with probability T_i^{n-1} , and will earn $v - (t_i - x) = x + \frac{1}{2} - T_i$ when she wins. So if, instead, i had deviated from the symmetric equilibrium by bidding a small amount ε more, as if she had signal $t_i + \varepsilon$, she would win $x + \frac{1}{2} - T_i(-\varepsilon)$ with additional probability $(T_i + \varepsilon)^{n-1} - (T_i)^{n-1} \approx (n-1)\varepsilon T_i^{n-2}$, for small ε , but pay an additional ε in the T_i^{n-1} cases in which she would have won anyway. In equilibrium i must be just indifferent about the small deviation so, since she knows only that T_i is uniformly distributed on $[0, 1]$,

$$\int_{T_i=0}^1 \left[(n-1)\varepsilon T_i^{n-2} \left(x + \frac{1}{2} - T_i \right) - \varepsilon T_i^{n-1} \right] dT_i = 0$$

(we are omitting terms in ε^2 and higher orders of ε)

$$\Rightarrow \left[\varepsilon T_i^{n-1} \left(x + \frac{1}{2} \right) - (n-1)\varepsilon \frac{T_i^n}{n} - \varepsilon \frac{T_i^n}{n} \right]_{T_i=0}^1 = 0 \quad \Rightarrow \quad x = \frac{1}{2}.$$

So i bids $t_i - \frac{1}{2}$, and the price is set by the bidder with the highest signal, $t_{(1)}$, which equals

$$v - \frac{1}{2} + \left(\frac{n}{n+1} \right)$$

on average. So the expected revenue from the auction is

$$v - \left(\frac{1}{n+1} \right).$$

These results confirm the Milgrom and Weber revenue rankings of the standard auctions.

Finally, since signals are affiliated an optimal auction can extract all the surplus for the auctioneer (see section 1.6 and Appendix 1.C). Here it suffices to ask each bidder to declare t_i , allocate the good to the high bidder (say) at the “fair” price $\frac{1}{2}[t_{(1)} + t_{(n)}]$, and ensure truth-telling behavior by imposing large fines on all the bidders if $t_{(1)} - t_{(n)} > 1$.

APPENDIX 1.E: BIBLIOGRAPHY

Sections 1.2–1.14 of this bibliography correspond to those sections of this chapter. Articles marked (*) are reproduced in *The Economic Theory of Auctions* (see section 1.18 below).

1.1 Survey and Guide to the Literature

Klemperer (*1999a).

1.2 Early Literature

Vickrey (*1961); Vickrey (*1962); Griesmer, Levitan, and Shubik (*1967); Ortega Reichert (*1968a); Wilson (*1969); Friedman (1956); Wilson (1967); Rothkopf (1969, see also 1980); Capen, Clapp, and Campbell (1971); Vickrey (1976).

1.3 Introduction to the Recent Literature

McAfee and McMillan (*1987a); Maskin and Riley (*1985); Riley (1989a).

1.4. The Basic Analysis of Optimal Auctions, Revenue Equivalence, and Marginal Revenues

Myerson (*1981); Riley and Samuelson (*1981); Bulow and Roberts (*1989); Vickrey (*1961); Vickrey (*1962); Bulow and Klemperer (*1996); Harris and Raviv (1981).

1.5. Risk-Aversion

Maskin and Riley (*1984); Matthews (*1987); Matthews (1983); Waehrer, Harstad, and Rothkopf (1998).

1.6. Correlation and Affiliation

Milgrom and Weber (*1982a); Crémer and McLean (*1985); Levin and Smith (*1996a); Crémer and McLean (1988); McAfee, McMillan and Reny (1989); McAfee and Reny (1992); Perry and Reny (1999).

1.7 Asymmetries

1.7.1 Private Value Differences

McAfee and McMillan (*1989); Maskin and Riley (*2000b, *1985); Griesmer, Levitan, and Shubik (*1967); Marshall, Meurer, Richard, and Stromquist (1994); Rothkopf, Harstad, and Fu (2003).

1.7.2 Almost-Common-Values

Bikhchandani (*1988); Klemperer (*1998); Bulow, Huang, and Klemperer (1999); Bulow and Klemperer (2002).

1.7.3 Information Advantages

Milgrom and Weber (*1982b); Milgrom (*1981); Engelbrecht-Wiggans, Milgrom, and Weber (1983).

1.8. Entry Costs and the Number of Bidders

1.8.1 Endogenous Entry of Bidders

Levin and Smith (*1994); Engelbrecht-Wiggans (*1993); Matthews (*1984); Fishman (*1988); Engelbrecht-Wiggans (1987); McAfee and McMillan (1987c); McAfee and McMillan (1988); Harstad (1990); Menezes and Monteiro (2000); Persico (2000b); Gilbert and Klemperer (2000).

1.8.2 The Value of Additional Bidders

Bulow and Klemperer (*1996).

1.8.3 Information Aggregation with Large Numbers of Bidders

Wilson (*1977); Milgrom (*1981); Milgrom (1979); Pesendorfer and Swinkels (1997).

1.8.4 Unknown Number of Bidders

Matthews (*1987); McAfee and McMillan (1987b); Harstad, Kagel, and Levin (1990); Levin and Smith (1996b); Piccione and Tan (1996).

1.9 Collusion

Robinson (*1985); McAfee and McMillan (*1992); Hendricks and Porter (*1989); Graham and Marshall (1987); Graham, Marshall, and Richard (1990); Mailath and Zemsky (1991); Hendricks, Porter, and Tan (1999).

1.10 Multi-unit Auctions

1.10.1 Optimal Auctions

Maskin and Riley (*1989); Palfrey (*1983); Avery and Hendershott (2000); Armstrong (2000); Rothkopf, Pekeč, and Harstad (1998).

1.10.2 Simultaneous Auctions

Wilson (*1979); Back and Zender (*1993); Anton and Yao (*1992); Klemperer and Meyer (*1989); Hansen (*1988); Maxwell (1983); Bernheim and Whinston (1986); Anton and Yao (1989); Daripa (1996a); Daripa (1996b); Nyborg (1997); Engelbrecht-Wiggans and Kahn (1998a); Engelbrecht-Wiggans and Kahn (1998b); Wang and Zender (2002).

1.10.3 Sequential Auctions

(i) *Bidders who demand only a single unit each*: Milgrom and Weber (*2000); Bulow and Klemperer (*1994); McAfee and Vincent (*1993); Bernhardt and Scoones (1994); Engelbrecht-Wiggans (1994); von der Fehr (1994);

Gale and Hausch (1994); Robert, Laffont, and Loisel (1994); Beggs and Graddy (1997); McAfee and Vincent (1997).

(ii) *Bidders with multi-unit demand*: Weber (*1983); Ortega Reichert (*1968b); Hausch (1986); Pitchik and Schotter (1988); Black and de Meza (1992); Krishna (1993); Robert (\approx 1995); Gale and Stegeman (2001); Pitchik (1995); Gale, Hausch, and Stegeman (2000); von der Fehr and Riis (1999).

1.10.4 Efficient Auctions

Ausubel (1998); Ausubel and Cramton (1998a,b); Dasgupta and Maskin (2000); Jehiel and Moldovanu (2001); Perry and Reny (1998); Bikhchandani (1999).

1.11 Royalties, Incentive Contracts, and Payments for Quality

Riley (*1988); Laffont and Tirole (*1987); Che (*1993); McAfee and McMillan (1986); McAfee and McMillan (1987d); Riordan and Sappington (1987); Branco (1997).

1.12 Double Auctions, etc.

1.12.1 Double Auctions

Chatterjee and Samuelson (*1983); Wilson (*1985); Rustichini, Satterthwaite, and Williams (*1994); McAfee (*1992); Leininger, Linhart, and Radner (1989); Satterthwaite and Williams (1989a); Satterthwaite and Williams (1989b).

1.12.2 Related Two-Sided Trading Mechanisms

Myerson and Satterthwaite (*1983); Cramton, Gibbons, and Klemperer (*1987).

1.13. Other Topics

1.13.1 Budget Constraints

Che and Gale (*1998); Pitchik and Schotter (1988); Pitchik (1995); Che and Gale (1996); Benoît and Krishna (2001).

1.13.2 Externalities between Bidders

Jehiel and Moldovanu (*1996); Jehiel, Moldovanu, and Stacchetti (1996); Caillaud and Jehiel (1998).

1.13.3 Jump Bidding

Avery (*1998); Fishman (*1988); Daniel and Hirshleifer (1995).

1.13.4 The War of Attrition

Bulow and Klemperer (*1999); Riley (1980); Bliss and Nalebuff (1984); Fudenberg and Tirole (1986); Krishna and Morgan (1997).

1.13.5 Competing Auctioneers

McAfee (*1993); Peters and Severinov (1997); Burguet and Sákovics (1999).

*1.14. Testing the Theory**1.14.1 Empirical*

Laffont (*1997); Hendricks and Porter (1988); Hendricks and Paarsch (1995); Laffont, Ossard, and Vuong (1995); Porter (1995); Laffont and Vuong (1996).

1.14.2 Experimental

Kagel (*1995).

*1.15. More on Specific Auction Forms**1.15.1 More on First Price Auctions*

Lebrun (1996); Maskin and Riley (2000a); Maskin and Riley (forthcoming); Athey (2001); Lizzeri and Persico (2000).

1.15.2 More on Second Price Auctions

Rothkopf, Teisberg, and Kahn (1990); Bikhchandani and Riley (1991); Bulow, Huang, and Klemperer (1995); Lopomo (1998); Wilson (1998); Bulow and Klemperer (2002).

1.16 Miscellaneous

Cassady (1967); Shubik (1983); Ashenfelter (1989); McAfee and McMillan (1994); McAfee and McMillan (1996); Riley and Li (1997); Ginsburgh (1998); Bulow and Klemperer (1998, Appendix B); Milgrom (2004).

1.17 Surveys

Klemperer (*1999a); McAfee and McMillan (*1987a); Milgrom (1985); Weber (1985); Milgrom (1987); Milgrom (1989); Riley (1989b); Maskin (1992); Wilson (1992); Bikhchandani and Huang (1993); Harstad and Rothkopf (1994); Rothkopf (1994); Wolfstetter (1996).

1.18 Collection of Articles

Klemperer (2000a).