An interesting feature of the telecommunications revolution and the growth of e-commerce over the past 20 years has been the explosion of auction markets. Millions of consumers all over the world now participate daily in auctions. In part, this reflects the fact that auctions are an exciting way to buy and sell. At least as important however has been the ability of the Internet to reduce significantly the costs of matching particular buyers and sellers. The leader in this development is of course e-Bay which through its operations in the U.S. and elsewhere auctioned over \$20 billion worth of goods among over 40 million confirmed users who either bid, bought or sold in 2005.¹

Auctions, however, started long before the Internet came along. Indeed, they have existed for thousands of years. Herodotus writes of a market for auctioning off wives in Babylonia in 500 BC. In AD 193, the Praetorian Guard auctioned off the rule of the Roman Empire to Marcus Didius Salvius Julianus for a bid of 25,000 sesterces. Today, the two most famous auction houses are probably Sotheby's (U.S.) and Christies (U.K.) each of which date back to the 1700s. These auction houses specialize in the sale of rare antiques and artwork, goods whose value is difficult to determine because there is a question of taste and opinion and because the value depends on market conditions that are difficult to forecast.

The sale of a good whose value is difficult to assess is exactly the kind of transaction that suits an auction because one needs a sizable number of interested buyers (or sellers) in order to get bids that reflect the spectrum of opinions as to the item's true value. Until recently, it has been much easier to incur the relatively high cost of bringing together such a critical mass of interested buyers for auctions focused on rather specialized markets, such as art, antique furniture, and race horses. Yet as we just noted, recent innovations in information technology and e-commerce have significantly lowered the cost of matching interested buyers and sellers. As a result, auctions are among the most popular sites and the fastest-growing business model on the Internet. Nor is the Internet the only venue for increasing auction use. Many markets that have traditionally been regulated such as electricity and telecommunications now use auctions for trading.

Auctions bring together different people with different values and different information. Whether one considers the potential buyers who go online to e-Bay or those who go to country

See Annual Report and the Silicon Valley Business Journal, various issues.

auctions, the fact is that the different bidders at any auction will typically have different valuations for the goods being sold at that site. This may simply reflect different preferences. For example, a decorator who prefers a colonial style will value a colonial rocking chair more highly than an Art Deco chair. The reverse is true for the Art Deco decorator. When each participant in an auction has a different or private value for the good being auctioned we say that the auction is a private value auction. Most often individual buyers' private values for the good will be known only to themselves.

Differences in information regarding the value of the good being auctioned, rather than differences in preferences, are another source of variation among auction participants. A common value auction is one in which the good being auctioned has one true value, but this value is not known to potential buyers. Moreover, the information that each potential buyer has for estimating the true value differs across the population of such buyers. A good example of a common value auction is the auction of rights to explore and drill for oil. In this case the bidders are oil companies who have performed tests on the oil tract and have made some estimate of the amount of oil present and then worked out a valuation for the rights to the oil tract. However, because the companies have performed the tests in different places they are likely to have different estimates and different valuations of how much the rights are worth. Yet there is just one true amount of oil present and thus only one true value of the rights to explore and drill for oil on the tract. The weekly bidding on the financial markets for U.S. Treasury securities offers another example of a common value auction.

For the most part in this chapter we will consider single unit auctions or those auctions in which a single unit of some good is to be auctioned. However, in many economic markets, such as in electricity or communications, there are multiple units of the good, which are being auctioned. These auctions are called multi-unit auctions.

25.1 A BRIEF TAXONOMY OF AUCTIONS

Before the advent of the Internet auctions often took place in a crowded room of hushed people. Those who wanted an item being auctioned off would raise their hands in response to the auctioneer's plea "Do I hear \$5, do I hear \$10 . . . ?" When the auctioneer finally cried "Going, going, gone" the auction ended and the last person to raise their hand won the good at the last price called by the auctioneer. This kind of auction is called an *ascending-bid* or *English* auction. It is probably the arrangement that the word auction first brings to mind for most people. However, the English auction is just one of a number of different auction types and not necessarily the most common. For example, rather than starting low and letting the price of the good rise with each successive bid, the auctioneer could have instead begun the auction with at an extremely high price—one that exceeds anyone's reasonable valuation. The auctioneer could then reduce the price and keep on reducing it until some one in the room raises their hand. This type of auction in which the first one to bid wins the good is called a *descending-bid* auction. It is in fact how flowers have long been sold at auction in Holland and for that reason is often called a *Dutch* auction.

Bidders participating in an ascending bid auction can watch the bidding and have multiple opportunities to place a bid. In a Dutch auction, bidders observe no bids other than the winning one. Once a price is reached at which someone will buy the item, the auction ends. In this respect, the Dutch auction is strategically identical to what is called a *sealed bid* auction. This is an auction in which the seller or auctioneer solicits a single bid in a sealed envelope from each interested buyer. The envelopes are then opened and the highest bidder

wins the auction. The similarity with a Dutch auction is that, here again, no bidder can observe any other bids.

Sealed-bid auctions, however, are not all the same. In particular, such auctions are usually divided into one of two types. One of these is referred to as a *first-price* auction while the second is referred to as a *second-price* auction. In a *first-price* auction the highest bidder pays the amount she bid whereas in a *second-price* auction she pays the amount bid by the *next* highest bidder.

Show that a dominant bidding strategy in an English auction is to continue bidding as long as the price in the auction is less than your true value of the good.

25.1

25.2 THE REVENUE EQUIVALENCE THEOREM

There is an interesting parallel between the four auction types that was first recognized by Nobel Laureate, William Vickrey in his classic 1961 paper. To see this, let's start by imagining that the chair of the economics department at your university decides to auction off a signed copy of the textbook you are currently reading at an English auction. Let's suppose that you would be willing to pay at most \$85.00 to win our book and that, unknown to you, this is the highest valuation of anyone in the class. Suppose further that the class is comprised of 170 students whose individual valuations may be ranked and which run from \$0.50 all the way up to your own value of \$85, increasing by 50¢ with each student. Bidding starts at \$0.50 and you and your classmates raise your hands as the chair bids up the price in 1¢ increments. With this procedure, it is inevitable that the price will eventually rise to the point at which only you and one other bidder remain, namely, at the price of \$84.50. At the next round however, when the chair increases the bid to \$84.51 your rival will drop out. You will then be the sole student with a hand raised. So you win the auction at a price of \$84.51.

Let's now consider what the outcome would have been if instead your chair had auctioned off the textbook using a second-price sealed bid auction. For the moment, let's assume that in this auction all the bidders write their true willingness to pay for the signed text on a piece of paper and put it in a sealed envelope. (We will show below that such a bidding strategy is in fact a dominant strategy in a *second-price* sealed bid auction.) With each student bidding her true reservation price, your bid will be \$85 and the next highest bid will be \$84.50. You will again be the winner. However, because it is a second-price auction, you will not pay \$85, but only \$84.50. Note that this is very close to the \$84.51 that you paid in the ascending bid auction. In short, if bidding one's true valuation is an optimal strategy, then the English auction and the second price sealed bid auction yield essentially identical outcomes.

Now let's investigate whether bidding based on one's true valuation is indeed optimal. To understand why it is, note that in a *second-price* sealed bid auction, the bid you submit only determines whether or not you win the auction. It does not affect the price that you will actually pay if you do win. That price is determined by the value of the second highest bidder. Therefore, bidding less than your true value of the good will only lower your chance of winning a *second-price* sealed bid auction. It does not change the price that you pay if you do win. There is then no advantage to bidding *less* than your true value. What

about bidding more? Increasing your bid above your true willingness to pay will increase the probability that you win only in the case when there is another bidder whose valuation is higher than yours. Otherwise increasing your bid has no effect. Yet if you win in these circumstances you will end up paying the other bidder's valuation—one that is higher than your own. In other words, you will end up paying more for the good than it is worth to you. Therefore bidding honestly is a dominant strategy in a *second-price* sealed-bid auction. As we have just seen, however, when everyone does this the outcome is equivalent to what occurs in an English auction.

Following Vickrey (1961), we have just established that an English auction and a secondprice sealed bid auction yield the same outcome if bidders pursue optimal strategies. This is important. For sellers who want to maximize their revenue from the auction, this fact implies that the choice between these two auction types is irrelevant. The seller will receive the same revenue either way. However, a seller may still wish to consider the two remaining types, the Dutch auction and the first-price sealed bid auction. What will happen in these two cases and how does this compare with the outcome of the English auction process? Here again, Vickrey (1961) provides the key insight.

We have already noted that the Dutch and first price sealed bid auctions are strategically identical. This is because they share two crucial features. One is that is in each case, a bidder has no additional information about the other bidders' valuations before making a bid. Instead, the bidder must simply bid based on her own valuation. The second feature common to these two auction types is that what one bids affects *both* one's chance of winning *and* what one pays. In both the Dutch auction and the first-price sealed-bid auction the winner pays the price that was bid to win the auction. This means that in each setting, bidders need to think strategically about what to bid in these settings.

Consider first a bidder's optimal strategy in a *first-price* sealed bid auction. Suppose that there are N bidders participating in the auction. Each bidder knows, of course, her own private valuation of the good being auctioned. Let us assume as well that each bidder knows the general distribution from which the other bidder's true valuations are draw. With these assumptions, it's easy to see that no bidder has an incentive to submit a bid above her true value. If a bidder did so and won the auction she would end up paying a price greater than her true value, which means she actually loses. We can also see that the bidder who has the highest valuation of the good should, if bidding optimally, win the auction. If the highest valuation bidder did not win because she submitted a bid that lost to say, the second highest bidder then she could always do better by increasing her bid slightly above the value of the bid that won and below her own.

These points suggest a possible winning strategy for any bidder. This is that she bid at least the valuation of the bidder with the next highest valuation. This insures that she will never lose to someone with a lower valuation of the good. The only difficulty is that, by assumption, no bidder knows the next highest valuation. Each only knows her own maximum willingness to pay. So how should a bidder proceed?

Clearly, a bidder needs to make an estimate of the next highest willingness to pay relative to her own. Let's consider how she might do this. Suppose that the bidder knows that all of the N valuations were drawn from a uniform distribution and denote her valuation by v. Because the bidder is particularly interested in the next highest valuation compared with her own she will focus on the distribution of valuations between 0 and v. This means that the bidder acts as though the remaining N - 1 bidders have valuations drawn from a uniform distribution over the interval [0, v]. The question then becomes what is the best guess of the next highest valuation among these remaining N - 1 bids? That is, what is the expected

value of the second highest valuation given that hers is the highest of N bids? Let's assume as we did before in the book auction that the values drawn by the other bidders are equally spaced on the interval. This assumption means basically that if we were to draw many, many samples of N - 1 values from the uniform distribution over [0, v], then the average value of the highest draw in these samples would be $\frac{N-1}{N}v$, while the average value of the second highest would be $\frac{N-2}{N}v$, and the average value of the third highest would be $\frac{N-3}{N}v$, and

so on. The lowest value on average would be $\frac{1}{N}v$.

For example, suppose that our bidder's valuation v is equal to \$85 and that there are 170 bidders in total participating in the auction. She can then proceed by assuming that the other bids below hers are equally spaced on the interval [0, 85]. The highest valuation in this interval would be $\frac{169}{170}$ \$85 = \$84.50, the next highest $\frac{168}{170}$ \$85 = \$84, the next to that $\frac{167}{170}$ \$85 = \$83.50, and so on. The bidder's optimal strategy in the first price sealed bid

auction therefore is to write down a bid of \$84.50. If she writes down more than this, she will pay more on average than is necessary. If she writes down less than this, she will on average lose the auction to someone who values the good less than she does.

The intuition of the foregoing argument is quite general. The bidder's objective is to acquire the auctioned good at the lowest possible price so long as that price does not exceed the bidder's own valuation of the object. For this reason, the bidder should condition her strategy on the assumption that her valuation is the highest because, if it is not, she will not wish to pay the price necessary to win. In turn, acting on the assumption that her valuation is the highest leads the bidder to bid the amount $\frac{N-1}{N}v$, which is the expected value of the

second highest bidder. Recall however, that the Dutch auction shares all the critical features of a first-price sealed bid auction and so is strategically equivalent to that case. Such strategic equivalence implies that the optimal strategy must be the same in each case. Hence, the optimal strategy in a Dutch auction is for the bidder to raise her hand as soon as the price N = 1

falls to $\frac{N-1}{N}v$. In short, the first-price sealed bid and the Dutch auctions both yield the

same outcome. Note also, that this is in fact yields the same winning bid as that which results in the English and second-price sealed bid auctions.

You are bidding for an original John Lennon hat in a sealed bid first-price auction. You are one of 8 bidders in this auction and the most you would be willing to pay for this hat is \$200. Show that your optimal strategy is to submit a bid of \$175.

In short, we have uncovered a very striking result. Regardless of whether the auction is English or Dutch, or first-price or second-price sealed bid, the outcome is the same. The winning bid in all four cases is identical. In our example, it is consistently \$84.50. This remarkable result is quite general and has been codified as auction theory's most well known

25.2

theorem, the Revenue Equivalence theorem.² Informally, the Revenue Equivalence theorem simply states that the expected revenue from an auction is the same regardless of which of the four basic types of auctions are used. A formal statement of Revenue Equivalence is given below.

Revenue Equivalence theorem (*private values*): Assume that there are N risk-neutral bidders each of which has a privately known valuation v of a good to be sold at auction with v drawn from a continuous distribution F(v) that is strictly increasing over the range $[\underline{v}, \overline{v}]$. Then any auction in which the object always goes to the buyer with the highest value of v, and in which any bidder with a value of \underline{v} enjoys an expected surplus of zero, results in exactly the same expected payment for each bidder v and yields exactly the same revenue to the seller.

The Revenue Equivalence theorem is a powerful result, in part, because it implies that auction design is not really an issue. There are, however, a number of conditions necessary for Revenue Equivalence to hold. For example, it must be the case that the auctioneer is understood by all bidders to report honestly the true value of the second-highest bid in a secondprice sealed bid auction (often called a Vickrey auction in honor of his pioneering work). The problem is that the actual bids tendered are known only to the auctioneer. Consequently, the auctioneer could increase the seller's revenue by pretending that a bid just under the maximum winning bid was received and declaring that this fictitious bid is the second highest price to be paid by the winner. Thus, if six bids of \$40, \$60, \$80, \$100, \$120, and \$140 are submitted, the auctioneer could report that the second-highest bid was actually \$139. The bidder offering \$140 will still win but pay \$19 above what she should have done. If this is a real possibility, then bidders in a Vickrey auction will reduce their bids in order to avoid such "rip-offs."

However, Lucking-Reily (2000) demonstrates that there may be ways to overcome the fear of auctioneer cheating in a Vickrey auction. Proxy bidding, the popular method of bidding on e-Bay, on an Internet auction is quite similar to a Vickrey auction. An online bidder submits a both an initial bid for the object as well as a maximum reservation bid. e-Bay then raises the bid incrementally on behalf of the bidder up to the stated maximum value— a value that is kept secret from other e-Bay users. For example, suppose that you submit an initial bid of \$20 for an out-of-print cover from a Grateful Dead album and that you simultaneously disclose to e-Bay that your maximum willingness to pay is \$100. If the bidding stops at a price of \$57, then that is the price that you will pay for the album cover. Of course, if this happens and if all bidders have, like you, submitted their true maximum willingness to pay then the reason that the bidding stops at \$57 is because that is the second highest valuation among all bidders. So, once again, the winning bid will be equal to the second highest valuation.

Of course, it is possible that \$57 was not the second highest price and that e-Bay is falsely claiming that value in order to claim a greater payment for itself. However, e-Bay makes considerable effort to persuade buyers that this is not the case. To begin with, e-Bay publishes a list of the losing bidders, their maximum bids and their email addresses after the auction closes. This permits the winning buyer to evaluate e-Bay's claim. In addition, e-Bay charges a relatively low commission from the sale of an object, roughly 5 percent or less.

² This result was first derived by Vickrey (1961) and then generalized by Myerson (1981) and Riley and Samuelson (1981).

Such a low commission means that e-Bay's primary interest is to encourage trades so as to garner as wide a circle of participants as possible. This too reduces their incentive to cheat. As noted, the great success of e-Bay strongly suggests that these tactics have worked and persuaded bidders to trust e-Bay's auctioneering. Properly designed then, the Vickrey auction should still be revenue equivalent to the others.

25.3 COMMON VALUE AUCTIONS

Suppose your university wants to offer coffee and light meals at the campus center and decides to auction off franchise rights to open a café. The auction type chosen is a first-price sealed bid design. Such an auction is called a common value auction because the café presumably has one true value, common to all participants. However, prior to the actual operation of the café that value is not known. Instead, each firm interested in bidding for the franchise can only estimate that value based on its own market research. Each will try to determine the expected number of students, faculty and staff that will eat there, what they will likely buy and at what price, and how much it will cost to serve them. However, it is crucial to remember that the true value of the café depends on what others are willing to pay for it. That is, although the bidders are different and will submit different bids, they are all trying to guess at the same thing, namely, the true market value of the café. What makes the bidders different here is not their personal valuation of the café but, instead, the fact that each has gathered somewhat different data or acquired different information regarding the café's true value. Presumably, if two bidders had collected exactly the same data they would have submitted identical bids. The usual case though is that each bidder will get a different signal and therefore form his or her own individual idea about the value of the café, even though ultimately that value is common to the entire market.

In general, the Revenue Equivalence theorem does not hold for common value auctions. This is why our formal statement of the theorem contains the parenthetical phrase (private values). The reason behind this non-equivalence in the case of common values is, however, quite subtle. In what follows, we try to illustrate the argument in an intuitive manner.

25.3.1 The "Winner's Curse"

Consider again our café example. After estimating the expected revenue and cost of operating the café, each firm interested in bidding will have an idea or estimate of what the franchise is worth and can use this estimate in submitting a bid. Of course, no two firms are likely to come up with exactly the same estimate of the café's value. Each firm is likely to survey different students or talk to different suppliers or otherwise base its estimate on information specific to that firm. Indeed, these differences in the information that each firm uses are a major source of the differences in the firms' estimates of the café's value and in the subsequent differences in their submitted bids.

After receiving all the tendered bids, the university will award the franchise to the highest bidder. The winner of the franchise will, of course, be the firm that made the highest bid. Yet in light of how the firms determined the amount of their bids, winning could in fact spell bad news. Because winning means that every other firm bid less for the franchise, it is quite possible that the winner overpaid for the franchise. This downside to winning is a central feature of common value auctions and is called the *winner's curse*. The curse is that the winner of a common value auction often turns out to be the loser because the winner

bids too much for the good. The franchise bidders will have collected different information about its true value and the winning bidder is likely the one who has the most optimistic information, and thus made the highest estimate of the value of the franchise. The winner's information is therefore the most likely to be upwardly biased. Consequently, a bid based on that information is likely to be too high.

Bidding in a common value auction therefore requires some sophistication. Let us continue to assume that the auction is a first-price sealed bid type. It should be clear that a bidder should not base his bid solely on the information he collected about the value of the good. He has to think about where this information came from, or alternatively, what kind of information the other bidders were likely to receive. Only by thinking this way can the bidder work out an estimate of the value of the good that will avoid the winner's curse. Suppose, for example, a bidder in our franchise auction knew that there were N bidders including him and that each bidder's estimate came from a uniform distribution, whose minimum is zero and whose maximum is \$50,000. The bidder could then work out that the mean value of this distribution is \$25,000, and that this value would, in fact, be the best or unbiased guess as to what is the true value of the franchise.

The difficulty is, of course, that a bidder is unlikely to know the upper limit or most optimistic estimate of the franchise's value. The bidder only knows the estimate he received and that the true value is uniformly distributed in the interval [0, U] where U is currently unknown to him. How should he proceed?

Consider a bidder whose research leads to an estimate of \$40,000 as the true vale of the café rights. However, he knows that this may well be an overestimate and he wants to avoid the winner's curse. One approach is for our bidder to assume that the estimate his research yielded is, in fact, the highest estimate obtained by any of the bidders. If so, then the bidder can use this information to work out a sensibly lower bid that should avoid the winner's curse.

What the bidder needs to do is to get a measure of the overall distribution of possible estimated values. The mean or average value of that distribution should be a good guess as to the true value of the café rights. He knows that the distribution runs from 0 to U. What he needs now is to get some idea of the value of U. Starting with the assumption that the \$40,000 is the highest of all the estimates drawn, let the bidder make our standard assumption that the estimates drawn by the other N - 1 bidders are uniformly distributed or equally spaced on the interval [0, U]. This assumption means that if we were to draw many, many samples of N values from the uniform distribution then the average value of the *highest* draw in these

samples would be $\frac{N}{N+1}U$, where U is the upper limit of the uniform distribution. Since

our bidder is assuming that his draw of \$40,000 is the highest draw he can work out U from N + 1

the equation $U = \frac{N+1}{N}$ \$40,000. For example, if the number of bidders were 10 in total

then U = \$44,000.

Why does the bidder assume that his own estimate is the largest? Since U is not known, the bidder recognizes that there is a positive probability that others may have drawn higher estimates. However, in determining his bid what the bidder really cares about is the true value of the café rights are *conditional* on his winning. So, the relevance of the higher estimates that others may have drawn is quite limited since if there are such higher estimates he is unlikely to win the café for himself. If our bidder does win the café though, then he can reasonably assume that his estimate was the highest or most optimistic. Since this is the scenario

in which the true value of the café is relevant, our bidder will build his bidding strategy around it, i.e., he will work on the assumption that his estimate is the highest and then work out the best bid that minimizes his winner's curse should he actually be the winning bidder.

Once our bidder has assigned some value to the upper limit U of the distribution of estimates, he can work out his best guess regarding the mean value of that distribution. Continuing with our assumption that the distribution of estimates is uniform over the interval [0, U] with U = \$44,000, the implied mean would be \$22,000. Accordingly this is our bidder's best estimate of the true value of the franchise. Note how much lower this bid is relative to the original value of \$44,000. This reduction in the estimated value of the café rights should therefore be quite effective in eliminating the winner's curse. If all the bidders in the common value auction calculate their bids in this way, each will shade his initial estimate in the same manner. As a result, it will still be the case that the bidder who initially drew the highest estimate will win the auction. However the "curse" of winning will be much reduced.

In short, the prospect of a winner's curse induces buyers in a common value auction to shade their bids below their individual estimate of an item's true value. The worse the winner's curse, the more such shading will occur. This is the intuition as to why, in a formal sense, the Revenue Equivalence Theorem does not generally hold for common value auctions.³ The reason is that different auction designs have different implications regarding the size of the winner's curse. In an English auction, for example, buyers get more and more information about the possible value as they watch the bid rise with each successive bid. In particular, they get increasing information about the lower bound of estimates for the café's value. In turn, this makes them more confident that the value is indeed high and so reduces the size of the winner's curse. Similarly, a second-price sealed bid auction can also lead to higher offers because the winner only pays the second highest bid. Of course, the weaker is the winner's curse effect the more aggressive will be the bidding and the greater the seller's revenue. The conventional ranking is that revenue is greatest for an English auction, next highest for Vickrey or second-price sealed bid auction and least for a first-price sealed bid auction which is equivalent to the case of a Dutch or descending price auction. However, this convention can break down when one considers slight departures from the standard common value case.⁴

Suppose your local town is auctioning off a franchise to sell hot dogs at the July 4th celebration. You and your partner decide to bid for the franchise. Including you there are eight groups bidding in the auction. Your market research on expected attendance, hot dog consumption and costs suggests that the franchise is worth \$2,000. Suppose you believe that your estimate, as well as the other bidders' estimates, are generated independently from a uniform distribution that starts at zero. What is your optimal bid for the franchise assuming that the distribution of values is uniform? If yours is the winning bid are you cursed?

25.3

³ Strictly speaking, it is not so much the common value aspect that generates the break from revenue equivalence as it is the fact that the ascending auction reveals to the ultimate winner information regarding the signals or estimates of those bidders who drop out in a manner which lets him use that information in setting his bids whereas he cannot use that information in a sealed bid process. Note also that Riley and Li (1997) show that the revenue difference between auction types may in practice be quite small, especially if the seller sets a sensible reserve price below which she will not sell.

⁴ This ranking originates from the famous paper of Milgrom and Weber (1982) on auctions with affiliated bidder values. See also, Milgrom (1989).

25.3.2 Almost Common Value Auctions

In our café example, we assumed that the true value was ultimately the same for everyone. The only difference among the bidders was the initial information that they had regarding precisely what that true value was. However, suppose that one of the bidders is the Starbucks chain and that winning the franchise is more valuable to it. This may be because winning will permit Starbucks to have a monopoly in the area, or because Starbucks can use its experience and buying power to operate the café more efficiently than can the other bidders. Whatever the reason, we will assume that if the café is truly worth v to all the other bidders, then it is worth v + \$1,000 to Starbucks. For example, if the true value of the café ultimately turns out to be \$25,000 for all the non-Starbucks buyers, it is worth \$26,000 to Starbucks. In this case, the café's value is not common to all buyers but the deviation from the common value case is relatively small. For this reason, this setting is often referred to as one of an almost-common value auction.

As it turns out, the small change involved in going from a common value auction to one with an almost common value can have very large consequences, as Klemperer (1998, 2002) in particular has emphasized.. The difference once again has to do with the winner's curse. In our example, all the non-Starbucks bidders face an exaggerated winner's curse. To beat Starbucks in the auction requires an extra large bid, but this just exacerbates the winner's curse. That extra \$1,000 that was driving Starbucks' bid will not be there for the non-Starbucks winner who has to outbid Starbucks. Bidders are not stupid, however. The non-Starbucks firms will recognize the exaggerated winner's curse that they face and therefore bid even more conservatively than in our earlier analysis. This permits Starbucks to bid more aggressively because it now faces a reduced winner's curse. Recognizing this makes the others bid less aggressively and so on. The end result is that Starbucks will always win the bidding but that it will do so at a much reduced price relative to that paid in a pure common value auction because of the extra conservative bidding pursued by its rivals. Hence, another reason that Revenue Equivalence may break down is because different auction designs may raise or lower the ability of dominant or advantaged buyers like Starbucks to exploit that uncommon but almost common advantage.⁵

25.4 AUCTION DESIGN: LESSONS FROM INDUSTRIAL ORGANIZATION

Auction markets have become increasingly common. Firms use auctions to purchase supplies, consumers bid for a variety of products through on-line auction sites, and the government auctions off Treasury bonds, mineral rights, and wireless spectrum licenses. In short, auction markets are common. It should come as good news therefore that the tools of industrial organization can yield insight into the operation of such markets. From a public policy perspective, such analysis is likely to be most useful in considering the auctions run by national and local governments. Ideally, such auctions will result in prices that are efficient in that they are close to the true market value of the item in question. In practice, this means that items will be auctioned to those for whom they have the highest (marginal) value. Since this

⁵ The "Wallet Game" introduced in Bulow and Klemperer (2002) is a very accessible introduction to the complexities of almost common value auctions

means that the revenue from the auction will be maximized, we can link efficiency to revenue maximization in considering auction design. Therefore, our evaluation of different auction mechanisms will largely be based on which auction design generates the most revenue. For cases in which the government is the seller as in say, the auctioning of airwaves, focusing on revenue maximization is appropriate for another reason as well. This is that the revenue raised can be used to reduce taxes and therefore to alleviate any tax-induced distortions.

Let us start by recognizing that government-run auctions are—except in cases such as auctions of obsolete military weapons-nearly always of the common value or almost common value type. Government bonds, mineral rights, and spectrum licenses are all items whose value to any one buyer depends largely on what other buyers would willingly pay for them. This makes life more difficult because the alternative case of purely private value auctions is certainly much simpler. The revenue equivalence theorem tells us that when bidders all have private values one type of auction is basically as good as another. Moreover, the auction outcome is efficient in that the winner will be the buyer who values the object the most and will pay a price equal to the valuation placed on the object by the buyer with the second-highest valuation, i.e., equal to the item's true opportunity cost as measured by the amount for which the winning bidder could sell it. Common value and almost common value auctions are, on the other hand, different. Here, revenue equivalence does not hold, and so auction design does matter. The source of revenue non-equivalence in this case can often be found in whether the auction design encourages bidders to enter the auction and in whether the design facilitates collusion among the bidders. Entry and collusion are of course issues that lie at the heart of industrial organization.

We begin by observing that our earlier analysis suggests that there is an advantage in an auction design that attracts many bidders. Recall that our optimum bid in the private value case was $\frac{N-1}{N}v = v - \frac{v}{N}$. Clearly, this bid increases as N grows.

However, it is worth noting that in common value auctions, the winner's curse also intensifies as the number of bidders who enter the auction increases. Recall our earlier café example only now assume that the auction is a Dutch or descending bid type. Suppose as before that you are again thinking of bidding \$22,000. If there are only 4 other bidders and none of them have indicated their willingness to buy as the price nears this point, you might not be too worried. If there are 40 other bidders, however, you might feel a good bit more hesitant. Being the high bidder out of a pool of 5 does not cause you to question very much your estimate of the café's value. Being the high bidder out of a pool of 41 though is quite different. When there are many bidders, the odds that your high estimate of the café's value is too generous rise considerably because the chances that someone will obtain an estimate far above the café's true value are much greater when there are 41 estimates than when there are 5. Thus, the more bidders there are the greater is the potential winner's curse and therefore, the more each buyer shades her bids. We generally expect an increase in the number of potential buyers to raise the price of a resource whose supply is fixed. Here however, the price-raising effect of more buyers is somewhat offset by the bid-reducing effect that enters into buyers' bidding strategies as they recognize the increasing winner's curse. Revenue maximization requires thinking carefully about which auction design can reduce the winner's curse.

As we noted before, the information revelation that accompanies an ascending auction can help to reduce the winner's curse in a common value auction. However, while an ascending auction reveals information and thus makes bidders more confident about their bid, it has the downside that the winner is unambiguously the one with the most optimistic

estimate and this may intensify the winner's curse. Further, while the ascending auction works well in the common value case, it is highly sensitive to the asymmetries that are present in the almost common value auction case. Recall what happened when one of the bidders for our hypothetical café was Starbucks. Because Starbucks valued the café at slightly more than everyone else this intensifies the winner's curse for all other buyers. In turn, this induces all other bidders to shade their bids and thereby reduce the winner's curse that Starbucks faces. Indeed, once other buyers realize Starbucks advantage, they may drop out altogether allowing Starbucks to win the café rights at a very low price. In other words, the information revelation on an ascending auction is almost perfectly designed to reveal Starbucks bidding advantage and thereby result in a low winning bid for Starbucks.

The foregoing point is illustrated by the 1995 ascending auction for mobile phone licenses in the Los Angeles area, a real world example discussed in Klemperer (2002). Pacific Telephone, the Baby Bell that then supplied fixed line telephone service, was certainly a well-known name in the area, and was widely reported to have made it clear to all potential bidders that it would bid whatever necessary to get the California market (see inset). While other firms did enter the auction for the Los Angeles license, the number of participants was much less than anticipated and the winning bid (paid by Pacific Telephone) of \$26 per head of population was much lower than had been predicted. Indeed, a similar auction in the Chicago area in which consumer demand for mobile phone service was almost certainly *less* than in Los Angeles, resulted in a noticeably higher winning bid of \$31 per head of population.

Another drawback to ascending auctions is that they may facilitate collusion. The repeated rounds of bidding do more than just reveal information about other buyers' estimated values. They also permit buyers to communicate and therefore to coordinate their bids. Cramton and Schwartz (1999) demonstrate this point in the context of the 1997 spectrum auctions conducted by the Federal Communications Commission (FCC). These auctions were ascending bid multi-unit auctions. A number of licenses were simultaneously auctioned off in an English auction. Cramton and Schwartz argue that different firms seemed to be signaling in their bids the identity number of the license areas that they were most interested in acquiring by matching the last three digits of their bid with their preferred area code, e.g., bidding \$313,378 for license area 378. Such a bid communicates to others that the bidder truly wants this area and will refrain from bidding aggressively for other regions so long as other bidders do not bid aggressively for area 378.

A similar story is told in Klemperer (2002) about a 1999 spectrum auction in Germany. In that case, there were ten licenses to be auctioned off and bids had to increase by a minimum of ten percent. One bidder, Mannesman, bid 18.18 million DM per MHz on licenses 1 to 5, and bid 20 million DM per MHz on licenses 6 to 10. Now observe that if you increase 18.18 by 10 percent the result is 20. T-Mobile, the other main bidder for these licenses, later admitted that it made just such a calculation. It concluded that Mannesman's bid represented something of an offer along the following lines: T-Mobile could have licenses 1 to 5 for 20 million DM per Mhz (the minimum amount it would need to beat Mannesman's bid) if it would not make any further bids for licenses 6 to 10. In fact, that is exactly what happened. The auction ended after just two rounds with all ten licenses going for 20 million DM per Mhz–well below anyone's estimate of the true willingness to pay of either T-Mobil or Mannesman.

Given the uncertain mix of advantages and disadvantages of an ascending auction, many have suggested the use of a first-price sealed bid design. In a sealed bid auction, the bidding process itself cannot be used to coordinate bids or to reveal asymmetries. Indeed, a sealed bid auction tends to encourage less advantaged bidders to participate. To understand why,

Reality Checkpoint This Wellcome Bid Could Have Been Higher

We have focused on the issues of auctions and auction design from the viewpoint of government auctions of scarce resources such as oil tracts or portions of the radio spectrum. However, the same principles apply for firms at the private level. Here again, the outcome of an auction depends crucially on such features as the number of bidders, the extent of any "winner's curse," and the nature of any asymmetries between bidders, among others.

For example, consider the merger of the two pharmaceutical firms Glaxo and Wellcome in 1995. In that year, Glaxo made an unsolicited bid for Wellcome worth \$14 million. At that time, the principal shareholder in Wellcome was The Wellcome Trust, a charity set up by the firm's founder, Sir Henry Wellcome, to finance medical research. Wellcome management was against the deal and sought out alternative buyers including two other drug firms, Zeneca and Roche. In response, Zeneca said it would be willing to bid \$15.5 million and Roche suggested it could enter a bid as high as \$17 million. However, each noted that it would be very costly to put together a viable bid-including the necessary financing-and so would only submit a bid if they were sure it would win. Glaxo publicly stated though that if another bid did come in, Glaxo would certainly top it. In the end, neither Roche nor Zeneca submitted bids. Without such rival bidders, Wellcome management could not persuade the trust to reject the offer. Glaxo acquire Wellcome for its initial bid of \$14 million. Judged by the terms of the drug mergers that quickly followed, e.g., Hoechst and Marion Merrell Dow, this price was about 15 percent below the norm. Glaxo's asymmetric willingness to pay acted to discourage entry by rivals and helped it win Wellcome at a bargain price, saving the firm hundreds of millions, if not over a billion dollars.

Source: R. Stevenson, "Wellcome Fails to Sway Trust on Sale," *New York Times*, January 28, 1995 (p. c1); and "Wellcome Cites Profit Gain in Move Against Glaxo Bid," *New York Times*, February 3, 1995.

consider a simple case with only two bidders. Suppose that both bidders know that each is drawing estimates of the item's worth from a uniform distribution ranging from \$0 to \$20. Bidder A has drawn a value of \$20. Bidder B has drawn a value of \$12. Neither of course knows what value the other has drawn. Clearly though, Bidder A will submit a bid rather well below her value of \$20. Bidding \$20 means winning with certainty but the net gain is zero. If she submits a bid of \$10, however, Bidder A will know that she still has at least a fifty percent chance of winning (if other players do not shade their bids) and enjoying a net of \$10. For Bidder B, the temptation to shade her offer is less compelling—especially if she thinks that a rival with a high value is going to underbid a lot. Thus, there is a chance for lower valued bidders to win a first-price sealed bid auction because unlike the ascending auction, the high-valued bidder has no recourse to outbid a rival in a subsequent round. In turn, precisely because a weaker bidder has some chance to win, a sealed bid auction can encourage more entry than an ascending bid auction.

The importance of attracting entry via a sealed bid process was potently illustrated by European mobile telephone auctions in 2000 and 2001. At that time, many governments were using auctions to allocate licenses to provide the so-called third generation (3G) of mobile telephone employing a new transmitting standard (Universal Mobile Telecommunications Service or UMTS). In the Netherlands, there were exactly five incumbent mobile telephone

companies and the government chose to auction off exactly five licenses using an ascending bid auction. It is easy to see the asymmetry between the five incumbent firms and any new bidders. The five incumbents clearly valued the new licenses more than an entrant. Thus the five incumbents were advantaged. Without additional entry, the fact that there were exactly five licenses being sold—one for each incumbent—reduced the auction to one that was very much like selling just one license to just one bidder in five separate cases. Yet the use of an ascending auction discouraged such additional participation (and may have permitted coordination between the five incumbents) with the result that only one additional and not very serious bidder emerged. In the end, the auction produced revenues of only 170 euros per capita. This was less than one-fourth the 650 euros per capita earned by a similar auction in the United Kingdom held just a bit earlier and less than one-third the predictions of the Dutch government. In contrast, a year later (after the dot.com and telecommunications booms had ended), the Danish government used a sealed-bid process to allocate four licenses to four incumbents and earned twice as much revenue as expected.

Since both an ascending (second-price sealed bid) and a descending (first price sealed bid) auction design have potential flaws, many have looked for ways to remedy these shortcomings. One suggestion is for the seller to announce a reserve price below which no bids will be accepted. To return to our earlier example, suppose that our seller knows that Bidder B's values are distributed between \$0 and \$10 while Bidder A's values are distributed between \$10 and \$20. If this information is known to the bidders as well, then in a first-priced sealed bid auction, Bidder A will never bid more than \$10 and may be tempted to bid a good bit less even when her own value is \$20. In an ascending English auction, Bidder A will always win but on average do so while paying only \$5. In either case, the seller would do well to introduce a reserve price of \$10 as she knows that at least one bidder will always value the item this much. Note how in this case, the reserve price mechanism preserves efficiency and insures that the item goes to the bidder with the highest value. Indeed, if the seller attaches a personal value to the item of above \$10, then she should choose that value as her reserve price to maintain an efficient auction outcome.

Another possible design alteration is to conduct a hybrid auction that combines both ascending and sealed-bid features. Klemperer (2002) for example, has suggested an Anglo-Dutch auction comprised of two rounds. The first round is an open ascending auction that ends when just two bidders are left. The second round is then run as a first-price sealed bid auction in which the reserve price is set as the final bid in the first round. The sealed bid procedure guards against collusion in the second round. It also encourages entry in the first round because, again, the sealed bid process permits a lower-valued bidder to win occasionally. At the same time, the first round process helps to establish a meaningful reserve price.

Summary

In this chapter, we have focused on the nature and implication of strategic interaction in the context of auctions. For private value auctions in which each bidder has her own valuation of the auctioned commodity, the four principal auction designs— English, Dutch, first-price sealed bid, and secondprice sealed bid—are revenue equivalent. That is, each yields the same winning bid and gives the commodity to the same winning bidder. For common and almost common value auctions, such revenue equivalence does not in general hold. In particular, for these cases, different auction designs can have very different outcomes both in terms of the revenue generated and regarding which bidder wins the auction.

Auctions have been increasingly used by governments to allocate scarce resources and license rights. Contests for ownership of firms and many other private battles can equally be viewed as auction processes. Because these are common or

almost common value auctions the link between auction design and auction outcomes is particularly important. In this regard, the key differences between auction designs reflect how each alternative works to modify the "winner's curse," encourage entry, and limit bidder collusion. Typically, an efficient auction will include a reserve price below which no bids will be accepted. However, setting that price can be difficult. A reserve price that is very low will impose little constraint on bidders' strategies. A price that is too high will deter many bidders. A starting point is for the seller to set a reserve price that is equal to her own valuation of the item.

Problems

- Consider Practice Problem 25.2 again in which you are bidding for an original John Lennon hat in a sealed bid first-price auction. In this case, however, let there be 20 other bidders in this auction. As before, assume that the most you would be willing to pay for this hat is \$200. Show that your optimal strategy is to submit a bid of \$190.
- 2. You are selling your house and want to get the highest price you can for it. What sort of auction would you prefer if:
 - a. you expect there to be 25 offers?
 - b. you expect the number of offers to be less than four?
- 3. When more than one buyer submits a bid for the same house, most states have strict laws forbidding real estate brokers to disclose to any one buyer the bid of any other buyer. Based on you're answers to 2a and 2b, what do you think is the justification behind such restrictions?
- At the time of the 2000 United Kingdom's auction of 3G telecommunication licenses, Britain had four incumbent mobile phone operators. Originally, it also planned to

References

- Bulow, J. and P. Klemperer. 2002. "Prices and the Winner's Curse." *Rand Journal of Economics* 33 (Spring): 1–21.
- Cramton, P. and J. A. Schwartz. 2000. "Collusive Bidding: Lessons from FCC Spectrum Auctions." *Journal of Regulatory Economics* 17 (May): 229–52.
- Klemperer, P. D. 1998. "Auctions with Almost Common Values: The 'Wallet Game' and its Applications." *European Economic Review* 42 (May): 757–69.

auction exactly four licenses. (In the end, it sold five.) Had it sold four licenses, the British planned to use a combined Anglo-Dutch approach. Under this design, the auction would proceed as an ascending auction until just five bidders were left. At that point, the auction would switch to a fourth-price sealed bid type in which the four licenses would be allocated to the top four bidders, each paying the price offered by the lowest successful bid. Comment briefly on this model. What economic considerations do you think were behind this auction design?

- 5. A house painter relates that most of his work is done for established customers for which he has little competition. He says that on these jobs, he submits a cost estimate and usually makes out reasonably well. Occasionally, however, the painter submits bids on jobs for less familiar customers for which business he is usually one of a number of painters the customer is considering. He says that for some reason, these jobs never work out so well and he always tends to lose money on them. Can you explain this result?
- —. 2002. "What Really Matters in Auction Design." *Journal of Economic Perspectives* 16 (Winter): 161–89.
- Lucking-Reily, David. 2000. "Vickrey Auctions in Practice: From Nineteenth Century Philately to Twenty-first Century E-Commerce." *Journal* of Economic Perspectives 14 (Summer): 183– 92.
- Milgrom, P. R. 1989. "Auctions and Bidding: A Primer." *Journal of Economic Perspectives* 3 (Summer): 3–22.

- Milgrom, P. R. and R. J. Weber. 1982. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50 (September): 1089–1122.
- Myerson, R. B. 1981. "Optimal Auction Design." *Mathematics of Operations Research* 6 (February): 58–73.
- Riley, J. G. and W. Samuelson. 1981. "Optimal Auctions." *American Economic Review* 71 (June): 381–92.
- Riley, J. G. and H. Li. 1997. "Auction Choice." Mimeo. University of California at Los Angeles.
- Vickrey, William. 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance* 16 (March): 8–37.