Chapter 1
Getting to Work

The era of putting auction theory to work began in 1993-94, with the design and operation of the radio spectrum auctions in the United States. Although the economic theory of auctions had its beginnings in the early 1960s, early research had little influence on practice. Since 1994, economic auction theorists have designed spectrum sales for countries on six continents, electric power auctions in the US and Europe, CO_2 abatement auctions, and various asset auctions. By 1996, auction theory had become so influential that its principal founder, William Vickrey, was awarded a Nobel Prize in economic science. In 2000, the US National Science Foundation's 50th anniversary celebration featured the success of the US spectrum auctions to justify its support for fundamental research in subjects like game theory. By the end of 2001, just seven years after the first of the large modern auctions, the theorists' designs had powered worldwide sales totaling more than $100 billion. The early US spectrum auctions had evolved into a world standard, with their major features expressed in all the new designs.

It would be hard to exaggerate how unlikely these developments seemed in 1993. Then, as now, the status of game theory within economics was a hotly debated topic. Auction theory, which generated its main predictions by treating auctions as "games," had inherited the controversy. At the 1985 World Congress of the Econometric Society, a wide gulf developed between bargaining theorists, who were skeptical that game theory could explain much about bargaining or be useful for improving bargaining protocols, and researchers in auctions and industrial organization, who believed that game theory was illuminating their fields. Although game theory gained increasing prominence throughout the 1980s and had begun to influence the leading graduate textbooks by the early 1990s, there was certainly no consensus about it in 1994, when the Federal Communications Commission conducted the first of the new spectrum auctions.

The traditional foundations of game theory incorporate stark assumptions about the rationality of the players and the accuracy of their expectations that are hard to reconcile with reality. Economic experimenters have tested the predictions of auction theory in laboratory experiments with human bidders and found many violations, but some key tendencies predicted by the theory do find experimental support. The findings indicate that existing theories oversimplify the way humans play games, and that real world auction design must be undertaken like other practical arts, by mixing theory with experiments and practical judgment.

Whatever the doubts in the academy, the dramatic case histories of the new auctions seized public attention. A 1995 New York Times article hailed one of the spectrum auctions as "The Greatest Auction Ever,"[1] and, in 2001, an academic auction designer—Professor Ken Binmore—was named a "Commander of the British Empire" for his work: the British 3-G spectrum auction[2] had raised more money than any other auction in history. In 1995, Professor Alvin Roth was called upon to apply game theory to revise the

---

[1]William Safire, "The Greatest Auction Ever," *New York Times*, March 16, 1995, page A17, commenting on FCC Auction #4.
[2] Designed by economists Ken Binmore and Paul Klemperer.

National Resident Matching Program—by which 20,000 US physicians are matched annually to hospital residency programs. The medical match, whose design is subtly and deeply connected to the new auction designs, was implemented in 1998.[3] By the mid-nineties, thirty-five years' of theoretical economic research concerning auction markets and matching markets was suddenly bearing very practical fruits.

Politics Sets the Stage

To most telecom industry commentators, the main significance of the US spectrum auctions was that a market mechanism was used at all. Spectrum rights (licenses) in the US and many other countries had long been assigned in *comparative hearings*, in which regulators compared proposals to decide which applicant would put the spectrum to its best use. The process was hardly objective: it involved lawyers and lobbyists arguing that their plans and clients were most deserving of a valuable-but-free government license.[4] With its formal procedures and appeals, a comparative hearing could take years to complete. By 1982, the need to allocate many licenses for cellular telephones in the US market had overwhelmed the regulatory apparatus, so Congress agreed to allow licenses to be assigned randomly among applicants by lottery.

The lottery did speed up the license approval process, but it created a new set of problems. Lottery winners were free to resell their licenses, encouraging thousands of new applicants to apply for licenses and randomly rewarding many with multimillion-dollar prizes. The lottery winners were often simple speculators with no experience in the industry and no intention of operating a telephone business. The hundreds of thousands of applications wasted economic resources on a huge scale, and the consequent need for the real wireless operators to negotiate and buy licenses from these speculators resulted in still more waste. The lotteries of small licenses contributed to the geographic fragmentation of the cellular industry, delaying the introduction of nationwide mobile telephone services in the United States.

A better process was needed, and Congress chose auctions as the answer. The question of how an auction market for radio spectrum should be designed and implemented called for fresh thinking and critical analysis.

Designing for Multiple Goals

When the US Congress authorized the first spectrum auctions as part of the 1993 budget act, it included several explicit instructions. One was that the first auctions were to be run in that fiscal year. A second called for wide participation in the new industry. The FCC initially responded to this second mandate by introducing bidding credits and favorable financing terms for small businesses and woman- and minority-controlled businesses, to reduce the cost of any licenses acquired by those businesses. Another goal of the auction was also to promote "efficient and intensive use" of the radio spectrum. The meaning of the word "efficient" was initially subject to debate, but it was eventually

---

[3] A similar match was implemented soon afterwards for clinical psychology fellowships, also with Roth's assistance.

[4] The process was once characterized by an FCC Commissioner as "the FCC's equivalent of the Medieval trial by ordeal." (See the dissenting statement of Commissioner Robinson in Re: Cowles Florida Broadcasting, Inc. et al, 60 FCC 2d 372 (1976)).

read in economic terms to mean, in the words of Vice President Albert Gore, "putting licenses into the hands of those who value them the most."[5]

There is a powerful tradition in economics claiming that markets, left to their own devices and operating in a sound legal framework, can achieve efficient allocations, but that tradition should not be applied too quickly to spectrum allocation. Even computing the efficient allocation can be an inhumanly hard problem, and getting participants to reveal the information necessary to do that computation is often impossible. Comparing the development of a universal standard (GSM) for mobile telephones in Europe the more fragmented system that emerged in the US highlights that the lottery system did not lead to efficient spectrum allocations. Getting the allocation nearly right the first time does matter. Achieving that with an auction system called for fresh thinking and critical analysis.

The actual task of designing and running the auctions in the United States fell to the Federal Communications Commission (FCC), which had no previous auction experience. Within the FCC, the auction design task was assigned to a group led by Dr. Evan Kwerel—an economist and long-time advocate of using auctions to allocate spectrum licenses.[6]

Like any other important FCC decision, the auction design decision would need to be based on an adequate public record—a requirement that would force the FCC to go through a long series of steps. It would need to write and issue a proposed rule, allow a period for comments and another for "reply comments," meet with interested parties to discuss and clarify the points of disagreement, resolve those disagreements, issue a ruling, consider appeals, and finally run the auction. Steps like these often stifle innovation, but that is not what happened on this occasion. With no political guidance about what kind of auction to use, no in-house experts lobbying to do things their way, and no telecom with an historically fixed position about how an auction should be run, Dr. Kwerel had unusual freedom to evaluate the alternatives. As matters evolved, FCC chairman Reed Hundt came to regard the auction as an opportunity to do something dramatic, novel and creative. The stage was set.

Kwerel drafted a notice that proposed a complex auction rule. Industry participants, stunned by the novel proposal and with little experience or expertise of their own, sought the advice of academic consultants. These consultants generated a flood of suggestions, and the FCC hired its own academic expert[7] to help them evaluate the proposed designs. In the end, the FCC adopted a kind of simultaneous ascending auction, in which bids increase over time for all licenses and bidders are free to switch their bids among licenses as information about the highest bids on various licenses emerges during the auction.[8]

---

[5] Quoted from Vice President Gore's speech at the beginning of FCC auction #4.
[6] Kwerel's initial advocacy can be found in Kwerel and Felker (1985), "Using Auctions to Select FCC Licensees," OPP working paper 16.
[7] Professor John McMillan.
[8] The final FCC rules resembled those included in two proposals submitted by economics professors. They most closely resembled the simultaneous ascending auction proposed by Professor Robert Wilson and me, including our simultaneous closings of bidding on all licenses and my *activity rule* to make that closing rule feasible. They also resembled the simultaneous ascending auction proposed by Professor Preston McAfee, which had no activity rule and proposed closing bidding on each license independently.

Substitutes and Complements

To understand the nature of the auction design problem, one must first understand the identities and needs of the bidders. In the initial PCS auction, there were three classes of potential bidders. The first group included long-distance companies with no existing wireless businesses. These companies, including MCI and Sprint, were making plans to enter the wireless business on a national scale. Each wished to acquire a license or licenses that would cover the entire United States, allowing it to make its service ubiquitous and to combine wireless with their own long distance service to offer an attractive and profitable package to consumers.

A second group comprised the existing wireless companies, including both giants like AT&T and some smaller companies. The companies in this group already owned or controlled licenses that enabled them to offer services to parts of the country. Their objectives in the auction were to acquire licenses that filled in the varying gaps in their existing coverage and perhaps also to expand to new regions or to the entire nation. These companies posed a regulatory challenge for the FCC, which wanted to allow them to meet their legitimate business needs without gaining control of so much spectrum that they could manipulate market prices. To avoid that outcome, the FCC had imposed limits on the amount of spectrum that any company could control in any geographic area. These wireless companies would be ineligible to bid for a nationwide PCS license of the sort that was typically awarded in European countries. From MCI's perspective, that meant that a nationwide license might be bought cheaply at auction, so it lobbied the FCC to structure the new licenses that way.

The last group consisted mainly of new entrants without wireless businesses. Some of these companies, like Pacific Bell in California, were quite large. These companies typically sought licenses or packages covering large regional markets, but not licenses covering the entire nation.

One of the first lessons to take from this description is that the auction game begins long before the auction itself. The scope and terms of spectrum licenses can be even more important than the auction rules for determining the allocation, because a license can directly serve the needs of some potential bidders while being useless to others. For the actual PCS auctions, a license provided its owner the right to transmit and receive radio signals suitable for mobile telephone service in a particular band of radio frequencies and in a particular geographic area. These license specifications constrain the possible spectrum allocations. The task of the auction designer was to promote the best (most "efficient") possible allocation, subject to those constraints.

Achieving efficiency involves various subtle complications. A certain license may be valuable to one bidder because it helps exclude entry and increase monopoly power, while it is valuable to another because the buyer will use it to create valuable services. In comparing the efficiency of allocations, only the second kind of value counts, but bidders don't respect that difference when placing their bids. The value of a license to a bidder may depend not only on the license itself but also on the identities of other licensees and the technologies they use, because that can affect their "roaming arrangements"—which allow their customers to use another company's services when they roam to the other company's license area. A third complication is that the bidders may need to pool information even to determine their own likely profits from various arrangements, for example because the bidders have different information about the available technology or forecasted demand.

But the fundamental barrier to efficiency that was most debated among the FCC auction designers concerned the "packaging problem." The value of a license to a bidder is not fixed; it generally depends on the other licenses the bidder receives. For example, a bidder might be willing to pay quite a lot for a package of, say, five or six licenses, but not much for smaller packages and not much extra for larger packages.[9] *Until such bidder knows all of the licenses it will have, it cannot say how much any particular license is worth.*

Consider a situation with just two licenses. If acquiring one license makes a bidder willing to pay less for the second, then the licenses are *substitutes*. If acquiring one makes the bidder willing to pay more for the second, then the licenses are *complements*. With more than two licenses, there are other important possibilities, and these add considerable complexity to the real auction problem. For example, if there are three licenses—say A, B and C—and a certain bidder anticipates needing exactly two of them to establish its business, then A and B are complements if the bidder has not acquired C, but they are substitutes in the bidder has already acquired C. Nevertheless, most economic discussions of the auction design are organized by emphasizing the two pure cases.

Recent auctions devised by economic theorists are most distinguished from their predecessors in the ways they deal with the problems of substitutes and complements. Our later analyses will show that some of the new designs deal effectively with cases in which the items to be traded are substitutes, but that all auctions perform significantly worse in the more general case in which licenses might sometimes be substitutes and sometimes complements. The impaired performance may take the form of loss of efficiency of the outcomes, uncompetitively low revenues to the seller, or vulnerability to collusion.

To illustrate how value interdependencies affect proper auction design, we turn to a case study in which the matter received too little attention.

New Zealand's Transponder Auction

New Zealand conducted its first auctions of rights to use radio spectrum in 1990. Some of the rights took the traditional form of "*license rights*" to use the spectrum to provide a specific service, such as the right to broadcast television signals using those frequencies. Others consisted of "*management rights*" according to which the buyer may decide how to use the spectrum, choosing, for example, between television broadcasts, wireless telephones, paging, or some other service. In theory, when management rights are sold, private interests have an incentive to allocate spectrum to its most profitable uses, but the problem of coordinating uses among licensees can also become more complex.

Acting on the advice of a certain consulting firm, the New Zealand government adopted a *second price sealed-tender auction* for its first four auction sales. According to William

[9] An instance of this sort arose in the Netherlands spectrum auction in 1998, in which most of the licenses were for small amounts of bandwidth. New entrants were expected to need five or six such licenses to achieve efficient scale and make entry worthwhile.

Vickrey's (1961) original description of the second price auction, its rules are these: Each bidder submits a sealed tender. Then, the license is awarded to the highest bidder for a price equal to the *second* highest bid, or the minimum price if only one qualifying bid is made. The auction gets its name from the fact that it is the second highest bid that determines the price.

The very idea of a second-price sealed-tender auction strikes many people as strange when first they hear about it, but on closer analysis, the auction is not strange at all. In fact, it implements a version of the ascending ("English") auction similar to one familiar at electronic auction sites like eBay.

In an ascending auction, if a bidder has a firm opinion about what the item is worth, then it can plan in advance how high to bid – an amount that we may call the bidder's *reservation value*. At sites like eBay, the bidder can report that value privately to the auctioneer, who will make place on its behalf, as if instructed to compete up to the specified price, but no higher. If everyone bids that way, then the outcome will be that competition ends when the price rises to the second highest reservation value, or thereabouts (with differences due to the minimum bid increment). In effect, if everyone adopts such a strategy, then the ascending auction is really just the same as a second price auction. In such an auction, the strategic considerations are easy: just set the reservation value to what the thing is worth. A bidder can't affect its price much anyway[10] and this bid wins whenever it should.

The second-price auction has two advantages over most other designs. First, it duplicates the outcome of the ascending bid auction without any need to assemble the bidders together or to have them hire separate agents. Second, it presents each bidder with a simple strategic bidding problem: each merely has to determine its reservation price and bid it. There is no need to make estimates of the number of other bidders or their values, since those have no bearing on a rational bidder's optimal bid.

The second-price auction has a simple extension to sales of multiple identical items, and it, too, can be motivated by considering ascending auctions. For example, suppose there is such an auction rule with seven identical items (perfect substitutes) for sale, to be awarded to seven different bidders in a single ascending bid competition. An analysis much like the preceding one leads to the conclusion that the items will be awarded to the seven bidders with the highest values for prices approximately equal to the eighth highest reservation price. To duplicate that with a sealed-tender auction, the rule must award items at a uniform price equal to the highest rejected bid. In such an auction, the right advice to bidders is simple: "bid the highest price you are willing to pay." A similar uniform-price rule has sometimes been used in the sale of U.S. Treasury bills.[11]

The New Zealand government was indeed selling essentially identical licenses to deliver television signals to New Zealand audiences. On the advice of its consultants, it did not adopt this "highest rejected bid" rule, but chose instead to conduct simultaneous second-price sealed-tender auctions for each license. New Zealand's second-price rules would

[10]The order of bids can affect the price, so a bidder with very precise expectations could, in principle, care about the timing of bids. In practice, such precise forecasts of other's bids are rarely available. We assume here that bid increments are negligibly small, so the price is literally equal to the second highest value and the "order effect" is negligible.
[11]The Treasury rule set a uniform price equal to the lowest accepted bid.

work well in one case only: when the values of the items were independent—neither substitutes nor complements. In the actual New Zealand auction, it would have been difficult to give bidders good advice. Should a bidder bid for only one license? If so, which one? If everyone else plans to bid for just one license and picks randomly, perhaps there will be some license that attracts no bids. Bidding a small amount for every license might then be a good strategy. On the other hand, if many spread around small bids like that, then bidding a moderate amount for a single license would have a high chance of success. When the values are so interdependent, unlinked, independent auctions inevitably involve guesswork that gets in the way of an efficient allocation.

Table 1

Winning Bidders on Nationwide UHF Lots 8 MHz License Rights			
Lot	Winning Bidder	High Bid (NZ$)	Second Bid (NZ$)
1	Sky Network TV	2,371,000	401,000
2	Sky Network TV	2,273,000	401,000
3	Sky Network TV	2,273,000	401,000
4	BCL	255,124	200,000
5	Sky Network TV	1,121,000	401,000
6	Totalisator Agency Board	401,000	100,000
7	United Christian Broadcast	685,200	401,000

Source: Hazlett (1998).

The actual outcome of the first New Zealand auction is shown in Table 1. Notice that one bidder, Sky Network TV, consistently bid and paid much more for its licenses than other bidders. Totalisator Agency Board, which bid NZ$401,000 for each of six licenses, acquired just one license at a price of NZ$100,000, while BCL, which bid NZ$255,000 for just one license, paid NZ$200,000 for it. Without knowing the exact values of various numbers of licenses to the bidders, it is impossible to be certain that the resulting license assignment is inefficient, but the outcome certainly confirms that the bidders could not guess one another's behavior. If Sky Network, BCL, or United Christian had been able to guess the pattern of prices, they would have changed the licenses on which they had bid. The bid data shows little connection between the demands expressed by the bidders, the numbers of licenses they acquired, and the prices they eventually paid, suggesting that the outcome was inefficient.

A second problem was even more embarrassing to New Zealand's government officials. John McMillan (1994) described it as follows: "In one extreme case, a firm that bid NZ$100,000 paid the second-highest bid of NZ$6. In another the high bid was NZ$7 million and the second bid was NZ$5,000." Total revenue, which consultants had projected to be NZ$250 million, was actually just NZ$36 million. The second-price rules allowed public observers to get a good estimate of the winning bidders' profits, some of which were many times higher than the price. To avoid further embarrassment, the government shifted from the second-price sealed tender format to a more standard "first-price" sealed-tender format, in which the highest bidder pays the amount of its own bid. As we will see later in this book, that did not guarantee in higher prices. It did, however, conceal the bidders' profits from a curious public.

The change in auction format still failed to address the most serious auction design problems. Unlinked auctions with several licenses for sale that may be substitutes or complements force a choice between the risk of acquiring too few licenses or too many, leaving a guessing game for bidders and a big role for luck. Allocations are unnecessarily random, causing licenses to be too rarely assigned to the bidders who value them the most.

Better Auction Designs

In the New Zealand case, alternative auction designs could have performed much better. For example, the government could have mimicked the design of the Dutch flower auctions. The winner at the first round would be allowed to take as many lots as it wished at the winning price. Once that was done, the right to choose next could be sold in the next auction round, and so on. With such an auction, no bidder would be forced to guess about which licenses to bid on. Each bidder could be sure that, if it wins at all, it will win the number of lots or licenses anticipated by its business plan at the bid price it chose.

There are other designs, as well, that limit the guesswork that bidders face. A common one in US on-line auctions allows bidders to specify both a price and a desired quantity. The highest bidders (or, in case of ties, those who bid earliest) get their orders filled in full, with only the last winning bidder running the risk of having to settle for a partial order. As with the Dutch design, efficiency is enhanced because bidders do not have to guess about which licenses to bid on, and such rules reduce the "exposure" risk that a bidder may wind up acquiring licenses at a loss, because it buys too few to build an efficiently scaled system.

The FCC Design and Its Progeny

In the circumstances of the FCC's big PCS auction, it was obvious that some licenses would be substitutes. For example, there would be two licenses available to provide PCS service to the San Francisco area. Since the two licenses had nearly identical technical characteristics and since, for antitrust reasons, no bidder would be allowed to acquire more than one, these licenses were necessarily substitutes. The argument that some licenses were complements was also made occasionally, but the force of the argument was reduced by the large geographic scope of the licenses.[12]

As in the New Zealand case, the main design issue was to minimize guesswork, allowing bidders to choose among substitute licenses based on their relative prices. When substitute goods are sold in sequence, either by sealed bids or in an ascending auction, a person bidding for the first item must guess what price it will have to pay later if it waits to buy the second, third, or fourth item instead. Incorrect guesses can allow bidders with relatively low values to win the first items, leading to an inefficient allocation. With this problem in mind, the final rules provided that the licenses would be sold all at once, in a single open ascending auction. The openness of the process would eliminate the guesswork, allowing bidders to switch among substitute licenses, and guaranteeing equal prices for perfect substitutes as well as an efficient outcome.

In order for the auction to work in this idealized way, bidding on all licenses would need to remain open until no new bids were received for any license. In a worst case scenario, the auction might drag on interminably as each bidder bid on just one license at a time, even when it was actually interested in eventually buying, say, 100 licenses. To mitigate this risk, the FCC adopted my suggestion of an "*activity rule*." In its simplest form, the rule prohibits any bidder from increasing the population covered by bids during the auction, as license prices increase, although it does allow bidders to switch back-and-forth among licenses. This rule ensures plenty of activity early in the auction and allows bidders to respond to changing prices, promoting a more efficient assignment of licenses.

The FCC rules have evolved slowly since the original 1994 design, but larger changes have been made to adapt the same ideas for other applications. One common variation arises when there are many units of each kind of item, such as power auctions with a few kinds of electricity contracts. In these auctions, the auctioneer accepts bids expressing total quantity demanded at a price; it then raises the prices slightly of goods for which demand exceeds supply. A series of "clocks" record the current prices for the various goods, and the rate of movement in the clock determines the progress of the auction. A similar clock auction was used in March, 2002 by the British government to buy 4 million metric tons of CO_2 emission reductions proposed by British businesses.

Clock auctions shares key characteristics with their FCC ancestor. Bidding on all items takes place simultaneously, so that bidders can respond to changing relative prices. Prices rise monotonically, ensuring that the auction progresses. All bids are serious and represent real commitments. There is an activity rule that prevents a buyer from increasing its overall demand as prices increase. Finally, bidding ends simultaneously on all the lots, so that opportunities for substitution do not disappear during the auction until all final prices are set.

[12] Dr. Mark Bykowsky of the National Telecommunications and Information Administration (NTIA) was a forceful advocate that licenses could be complements and proposed a complex package auction design to accommodate the possibility. His case that complementarity was important is more convincing for the later auctions in which smaller licenses were sold. Whatever the intellectual merits of this position, the short time available to run the first auction led to a near-consensus that the package auction proposal involved too many unspecified details and unresolved uncertainties to evaluate and adopt immediately.

New variations based on the same principles continue to be created to solve a wide range of economic problems. Electricité de France (EDF) used a particularly interesting one in 2001 in a sale of electrical power contracts. The sale involved power contracts of different lengths, ranging from three months to two years, but all beginning in January of 2002. Because different buyers wanted different mixes of contract lengths and because all contracts covered the first quarter of 2002, the seller needed to compare several different kinds of offers to determine which to accept.

Professors Larry Ausubel and Peter Cramton developed the auction design. The first step was to assist EDF in developing a standard for comparing bids on contracts of different lengths. During the auction itself, the auctioneer raised the prices of the various acceptable contracts using "price clocks," maintaining price differences that would leave the seller indifferent among the offers. It stopped raising prices when the total remaining demand exhausted the total power offered for the first quarter of 2002. Such an auction creates competition among bidders for contracts of different lengths, increasing both efficiency and sales revenue compared to more traditional auction designs.

Comparing Seller Revenues

The question most frequently asked of auction designers is this: What kind of auction leads to the highest prices for the seller? The answer, of course, must always be heavily qualified, but it still holds a surprise for many people. There is no systematic advantage of either sealed bids over open bid auctions, or the reverse.

A particular formal statement of this conclusion is known as the *payoff equivalence theorem*. It holds that in an important class of idealized situations, the average revenues from an auction and the payoffs of bidders are exactly the same. To illustrate the logic of the idea, suppose you are selling an item that is worth $10 to bidder A and $15 to bidder B. If you sell the item using an ascending bid auction with both bidders in attendance, then bidder A will stop bidding at a price close to $10 and B will acquire the item for that price. If you use sealed bids instead and sell the item to the highest bidder, then the outcome will depend on what the bidders know when they bid. If they know all the values, that in theory B will bid just enough to ensure that it wins—around $10 or $10.01 and A will likely bid just under $10, so the price will be just the same. In both kinds of idealized auctions, the seller receives about the same price in both cases and each party earns the same profit. As William Vickrey first observed, a similar conclusion holds on average both for a wider class of auction rules and in a wider class of situations than the one described here.

Practical people tend to feel puzzled when presented with Vickrey's irrelevance conclusion. Auctioneers who conduct ascending auctions often say that they generate more excitement and more competition than sealed bids. After all, they argue, no bidder is willing to bid close to its value unless pushed to do so by the open competition of the ascending auction design. Those who favor sealed-tenders counter by arguing that ascending auctions never result in more being paid than is absolutely necessary to win the auction, while sealed tenders sometimes leave very large sums of money "on the table." In the December 1997 auction for licenses to provide wireless telephone services in Brazil, an international consortium including Bellsouth and Splice do Brazil bid $2.45

billion in that auction to win the license covering the Sao Paulo concession. This was about 60% higher than the second highest bid, leaving nearly $1 billion on the table.[13]

Similar price debates have arisen in discussions of the rules used to sell Treasury bills in the United States. The Treasury staff have periodically argued the relative merits of two alternative auction schemes—one in which each bidder pays the amount of its own bid for each bill it buys and another in which all bidders pay the same price: the lowest acceptable bid or "market-clearing price." Advocates of the first ("each-pays-its-own-bid") scheme say that the government will get more money from the auction, since winning bidders are by definition people who have bid more than the lowest acceptable bid. Advocates of the second ("each-pays-the-market-clearing-price") scheme counter that bidders who know they must pay their own bid when they win will naturally bid less, reducing market clearing price and leading to lower revenues.

Informal arguments like the ones just described lead to no useful conclusions. A formal analysis based on the *payoff equivalence theorem* introduced in this book helps to cut through the confusion. We will find that if the allocation of lots among bidders is the same for the two designs, then the average payoffs to all parties, including the average prices obtained by the seller, must also be *exactly the same*. One cannot conduct a meaningful analysis of average prices alone, without also studying how the designs affect the distribution of the lots among the winning bidders.

When the assumptions of the *payoff equivalence theorem* reasonably approximate reality, the auction designer should shift its attention from how payments are determined to such other factors as the costs of running and bidding in the auction, timing the auction and packaging lots to attract bidders, the vulnerability of the auction to collusion among bidders or to corrupt behavior by the auctioneer, and so on. When the assumptions fail, something valuable is still gained from the theorem: attention is shifted to how the differences between the assumptions and the reality may make one auction form more effective than another.[14]

The Academic Critics

Economists who work at putting auction theory to work encounter a dazzling array of issues, from ideological to theoretical to practical. Recognizing the complexity of the problems and the short times available to solve them, the engineering work for auctions sometimes entails guesses and judgments that cannot be fully grounded in a complete economic analysis. Auction designers use theory to generate ideas, test the ideas when they can, and implement them with awareness of their limitations, supplementing the economic analysis with worst-case analyses and other similar exercises.

---

[13] While the 60% overbid may be atypical, the ordinary amounts of money left on the table are still impressive. For example, in the Brazilian band A privatization, the median overbid was 27%. That is, for half the licenses, the winning bidders bid *at least* 27% more than the second highest bid.

[14] This use of the revenue equivalence theorem is similar to the best uses of other important theorems in economic theory. In practice, the first welfare theorem, the Coase theorem, the Miller-Modigliani theorems, and monetary neutrality theorems are best used as starting points for an analysis. One uses these theorems to identify and reject plausible-sounding but incorrect arguments and to focus the analysis on how particular failed assumptions might alter the conclusion and guide the policy decision.

The idea that economic theorists can add value through this mixture of auction theory and practical judgment has come under attack from some members of the economics profession. Some of the more frequent attacks, and my responses to them, are expressed below.

Resale and the Coase Theorem

One of the most frequent and misguided criticisms of modern auction design comes in the form of the remarkable claim that the auction design doesn't matter at all. After all, say the critics, once the licenses are issued, parties will naturally buy, sell and swap them to correct any inefficiencies in the initial allocation. Regardless of how license rights are distributed initially, the final allocation of rights will take care of itself. Some critics went even farther, arguing on this basis that the only proper objective of the government is to raise as much money as possible in the sale, since it shouldn't and can't control the final allocation anyway.

To justify this argument, the critics relied on the Coase Theorem, which holds that if there are no frictions in the market and no wealth effects on preferences, then the initial allocation of property rights cannot affect the distribution of wealth in society. It cannot affect the efficiency of the allocation or anything that is relevant for productive efficiency. Coase reasoned that so long as the allocation remains inefficient, the parties will find it in their interests to buy, sell and swap as necessary to eliminate the inefficiency.

Whatever merits the Coasian argument may have in other situations, it plainly leads to the wrong conclusion in this case.[15] Auction and bargaining theory and the history of cellular telephones in the US teach us that the initial assignment of right does affect efficiency. The theoretical argument juxtaposes two well-known propositions. First, there exist auction mechanisms that can achieve efficient license rights allocations, even when there are many available licenses, provided the government uses the auction from the start. Second, even in the simplest case with just a single license for sale, there exists *no* mechanism that will reliably untangle an initial misallocation. The difference between bargaining and auctions is that in bargaining, parties will be inclined to exaggerate their position to gain a bargaining advantage. That unavoidable exaggeration often delays and sometimes blocks a mutually profitable agreement. In contrast, a simple English auction leads to an efficient allocation with a single item, and the generalized Vickrey auction extends that outcome to any number of licenses.[16]

In the actual situation in the United States, the bargaining problems among multiple parties were much harder than even the theory acknowledges, so the Coasian reasoning based on assuming that bargainers reach efficient agreements does not apply. The much slower development the cellular telephone industry in the United States than in Europe

[15] The Coase theorem has includes a variety of assumptions that may fail in this application, such as the assumption that the parties values reflect social value—not market power, the assumption that the parties have unlimited budgets, so spending on spectrum rights does not impair the ability to invest in infrastructure, and the assumption that rights have no "externalities," that is, that bidders don't care about which competitors get license rights. The importance of the last assumption is analyzed by Jehiel and Moldovanu (2001).

[16] This argument is developed more fully in chapter 3, after the relevant theory has been introduced.

demonstrates the importance of the initial steps. Consumers long ago demonstrated their willingness to pay amply for the ability to "roam" and use their mobile telephones across the nation, but US consumers today still face unnecessary gaps in coverage resulting from the industry's initial fragmentation.

Mechanism Design Theory

A second line of criticism emerges from a part of game theory called "*mechanism design theory*." A "mechanism" is essentially a set of rules to govern the interactions of the parties. For example, it may specify the rules of an auction. Are there to be sealed or ascending bids? If sealed bids, how will the winner and price be determined? And so on.

Once the rules of the mechanism and the designer's objective have all been specified, the designer applies some criterion or "*solution concept*" to predict the outcome and then evaluates the outcome according to the objective. In this highly mathematical theory, the ultimate aim is to maximize the performance according to the specified objective. For example, one might try to find the auction that maximizes the expected selling price or the expected efficiency of the outcome. We will treat parts of this theory at length later in this book.

Mechanism design theory poses this challenge to practical auction designers: how can you justify any use of theory without applying the mechanism design approach? If you believe your theory describes the behavior of players, why don't you use the theory to optimize the mechanism performance?

The answer is the same one applied to any far-reaching use of an optimization model. Optimization requires that one trust a model to be complete and accurate and the objective clear and fully specified. These extreme conditions are unlikely to be satisfied in unique and complicated situations. Yet even when a model is not complete, it can lead to insights that are useful to the designer. Just as a mechanical engineer whose mathematical model assumes a frictionless surface treats those calculations as inexact, an economic designer who assumes that the players are optimizers and have rational expectations may do the same. Just as the real-world mechanical engineer pays attention to factors that increase friction and builds in redundancy and safety margins, the real-world mechanism designer pays attention to timing and bidder interfaces to make rational decisions easier, and plans to accommodate worst-case scenarios, in case a few bidders behave contrary to expectations.

At the present state of the art, academic mechanism design theory relies on extreme assumptions to reach theoretical conclusions that can sometimes be very fragile. In the standard optimizing approach, the mechanism designers apply a game theoretic solution that assumes that bidders maximize accurately and are completely confident that others will maximize accurately as well. The bidders are assumed to hold this confidence unshakably, even when the consequences of a small mistake by another bidder can have dire consequences.[17] Of course, practical mechanism designers also use equilibrium theory, but a good analysis never stops at that. Mechanisms proposed for practice need to

---

[17] Even the most distinguished mechanism design theorists argue this way, and they publish their results in the most prestigious academic journals. For example, see Dasgupta and Maskin (2000) and Perry and Reny (2002).

be tested for robustness. Those that are too fragile should be rejected, while robust mechanisms should often be adopted even if they have no provable optimality properties.

Besides the very demanding behavioral assumptions that characterize the theoretical mechanism design approach, the formal models of the theory typically capture form only a small subset of the issues that a real auctioneer faces. Some of the important issues that are usually omitted from mechanism design models are listed below. While none of these is incompatible with mechanism design theory in principle, accounting for all in a single optimization model is beyond the reach of present practice.

- *What to sell?* If a farmer dies, should the entire farm be sold as a unit? Or should some fields be sold to neighbors? The house and barn as a holiday and weekend home? How should the FCC cut up the radio spectrum? Should power suppliers be required to bundle regulation services, or should that be priced separately?
- *To whom and when?* Marketing a sale is often the biggest factor in its success. Competitors, too, may try to discourage one another, in order to get a better price.[18] Auctioneers may seek expressions of interest in order to determine which bidders are best qualified to bid.
- *How?* For example, if the deal is complicated and needs to be individually tailored for each bidder, a seller might prefer to engage in a sequence of negotiations to economize on costs. If an auction is to be used, the right kind can depend, as we have already seen, on whether the items are substitutes or complements.
- *Interactions?* These decisions are not generally made independently. The desirability of selling the farmhouse separately depends on answering "to whom," that is, on the identity of the potentially interested buyers. And, the auction design may depend on whether there is potential competition between a buyer of the whole property and buyers of the parts.
- *Mergers and Collusion?* The European spectrum auctions of 2000, with their very high stakes, provided some interesting examples of before-the-auction actions to reduce competition. In Switzerland, last minute mergers among potential bidders resulted in only four bidders showing up for four spectrum licenses. The auction was postponed, but the licenses were eventually sold for prices close to the government-set minimum. Similar problems of valuable spectrum attracting few bidders and resulting in prices near the minimum occurred in Germany, Italy, and Israel.
- *Resale?* Most of the theory of mechanism design starts with a given set of bidders who keep whatever they buy. The possibility of resale not only affects auction strategy, it may also attract speculators who buy with the intention of reselling. Should the seller encourage speculators, as additional bidders create more competition in the auction? Or should it discourage them, since value

[18] On the eve of the FCC PCS spectrum auction #4, the author made a television appearance on behalf of Pacific Bell telephone, announcing a commitment to win the Los Angeles telephone license, and successfully discouraging most potential competitors from even trying to bid for that license.

captured by speculators must come from someone else's payoff—possibly the seller's?

The mechanism design extremist's view, which holds that the only consistent approach is to develop theoretically "optimal" mechanisms, is not useful in practice. Even if we could incorporate all the features described above, our models of human behavior are not nearly accurate enough for use in optimization. Behavior is neither perfectly stable over time, nor the same across individuals, nor predictable for any single individual. Useful analyses have to be cognizant of that.

Despite these limits, a large portion of this book focuses on mechanism design and related analyses. The theory is useful in practice for thinking through some issues and guiding some decisions. Among the decisions that the theory illuminates are ones about *information policy* (what information to reveal to bidders), how to structure *split awards* (in which a buyer running a procurement auction splits its business between two or more suppliers), how to create *scoring rules* (in which bids are evaluated on dimensions besides price), and when and how to implement *handicapping* (in which the auctioneer treats bids unequally in order to encourage more effective competition). The mechanism design approach also helps answer important questions about when to use auctions at all. Members of the Institute for Supply Management discuss this issue so often that they have coined the term "*auctionable*" to describe goods and services that can be most effectively purchased using auctions.

Theory and Experiment

In sharp contrast to mechanism design purists, some economic experimenters raise an opposite objection: why should any attention be paid to auction theory at all, now that we have the capability to test alternative auction designs in experimental economics laboratories? Theories sometimes fail badly. The rest of the time, they explain only some of the data, so why rely on theory at all?

The possibility of experimental tests has, indeed, fundamentally shifted the way auctions can be designed. In the FCC auction design, successful tests conducted by Professor Charles Plott at Caltech helped convince the FCC to adopt the theoretically motivated design. Working software demonstrating the design was another important element.[19] Yet, the experiments to date have been very far from replicating the actual circumstances of high value auctions.

In practice, it is unlikely that anyone will ever test a range of actual proposals in a completely realistic setting. The amounts at stake in experiments are necessarily much smaller, and the preparation time for bidders will normally be much less. Because experimental settings differ so much from the auctions they simulate, the role of theory is indispensable. Theory guides the design of experiments, suggests which parts of any

---

[19] Working software demonstrating the feasibility of the new design was another important element. Implementation issues also played a huge role in the debate. The very possibility of running the computer implemented simultaneous auction drew hackles from critics in 1994. To rebut the critics, my assistant, Zoran Crnja, programmed a flawless small-scale version of the software in a set of linked Excel spreadsheets. His software convinced the FCC that a reliable system could be created using our proposed rules even in the short time available.

experimental results might be generalized, and illuminates the economic principles at work, enabling further predictions and improvements upon the original design.

Lord Alfred North Whitehead, when asked whether theory or facts was more important, answered famously: “theory about facts.” Indeed, theories that are incompatible with facts are useless, but there can be no experimental designs and, indeed, no reporting of experimental results without a conceptualization of the issues. Theory will always play a key role in answering engineering questions, including questions about auction design.

Practical Concerns

The final criticism is that, in the real world, the auction rules are a secondary concern in setting up and running a complex auction. Several other issues are said to be more important.

One such issue is marketing: an auction cannot succeed without participants. A partial answer to this is that, depending on the circumstances, changing the auction rules may attract more participants. There are many examples of auctions and other competitions that get poor results because the rules are rigged to favor particular bidders, discouraging others from participating. The earlier description of MCI’s attempts to rig the US spectrum auctions in its favor by making the “lot” a single national license is one among many examples. When different bidders want different lots, a package design, such as the ones often used in bankruptcy sales, may enable wider participation.

A second important practical issue concerns the property rights being allocated. For example, if auctions are to be used to allocate take-off and landing rights at a congested airport, then the rights themselves need to be carefully defined. What is to happen to a plane that is delayed for mechanical reasons and cannot depart in its assigned slot? What about weather delays that decrease the capacity of the airport. It is certainly true that no sophisticated auction rule can lead to a good outcome unless this practical issue is resolved, but it is equally true that an auction system that fails to coordinate all the resources needed by the airlines—takeoff slots, landing slots, rights through *en route* choke points, gate access, and so on—cannot succeed regardless of how well rights are defined. Real problems require comprehensive solutions, and the auction rules are often an important part of that.

Yet another important practical detail for electronic auctions is the interface used by bidders. The original FCC auction software made it easy for bidders to make mistakes. On several occasions, bidders made what came to be called “fat finger bids.” For example, when trying to bid $1,000,000, a bidder might accidentally enter a bid of $10,000,000—an error encouraged by the fact that the early interfaces could not accept commas in the bid field.

The FCC’s solution for this problem, however, was one that considered more than the ease of bidding. Under the FCC’s initial rules, bidders found it easy to communicate messages, including threats, with their bids in the auction. Suppose, for example, that bidder A wishes to discourage competitor B from bidding on a particular license, say #147, in a particular auction. If B bids on that license, A might retaliate by raising the price of another license on which B has the current high bid of, say, $9,000,000 by bidding $10,000,147, where the last three digits send a none-too-subtle message about its motivations. Such bids were frequently observed in some of the early FCC auctions.

Both the "fat finger" and the signaling problems were solved when the FCC changed the auction interface to require that a bidder select its bid from a short drop-down menu on its bidding screen. All bids on the menu used round numbers, being the minimum bid plus one or more increments. This eliminated typos involving one or more extra digits and simultaneously made it much harder for bidders to encode messages in their bids.

Some critics respond to such anecdotes with the claim that while these do show that rules matter, they mainly show the dangers in electronic auctions or auctions using novel rules. However, even familiar, low-tech auctions can perform badly on account of problematic rules. In 1998, the Cook County, Illinois, tax collector conducted a traditional oral outcry auction to sell the right to collect certain 1996 property taxes that were two years overdue. In that "1996 tax sale" auction, a bid specified the *penalty rate* that the winning bidder could add to the taxes due, as compensation for its collection services. The auction was conducted in an ordinary meeting room, with the auctioneer sitting in the front. The auctioneer would read a property number and the bidding instantly began with the bidders shouting penalty amounts. The maximum opening bid was 18% and successively lower bids are shouted until a winning low bidder is determined.

The trouble occurred when several bidders simultaneously opened with bids of the maximum amount. Under the Cook County rules for that year, in the event of such a tie, the auctioneer was to assign the properties to winning bidders essentially at random. A bidder tied with, say, five others at 18% then faces a simple choice. It can bid less than 18%, having about a one in six chance to win the auction at a much lower rate than 18%. Or, it can sit quietly, having a one in six chance to win at a rate of 18%. Most bidders chose to sit quietly, and about 80% of the properties sold at the maximum rate of 18%.

How can we be sure it was the faulty rules, rather than collusion among (more than a dozen) bidders, that accounted for this outcome? A few days after the auction began, the county auctioneer announced a change in the rules. In the future, a tie bid at 18% would result in withdrawal of the property from the auction. After the change, penalty rates quickly collapsed to a lower level, leading some bidders to seek a court order restraining the auctioneer from changing its rules during the auction. The order was issued and the winning bids immediately returned to 18%.

Understanding auction theory is helpful for more than just avoiding obviously bad designs. Well-designed auctions that link the allocation of related resources can perform very much better than traditional auction sales. In the New Zealand case described earlier, even if the novel second-price auction rules had been replaced with more traditional pay-as-bid rules, the fact that the TV licenses were good substitutes means that any simultaneous sealed-bid auction is prone to lead to a misallocation. Computational experiments suggest that 25-50% of the value might have been lost simply because the allocation was so poorly coordinated. In similar circumstances, the simultaneous ascending auction design that is the current world-standard for spectrum auctions can theoretically lead to an efficient or nearly efficient outcome.

The simultaneous ascending auction has limits, too, which can be particularly important when the items for sale are ones that different bidders prefer to package in different ways, or when there are complicated constraints on the collection of acceptable offers. In such cases, a package auction design can both attract a wider set of bidders and vastly increase the likelihood that the right packages emerge from the auction. The design of these auctions is subject to many pitfalls, to which we return in section II of this book.

There are many more examples of painful lessons about the importance of the detailed auction rules. In the US electricity markets, ill-considered auction rules frequently contributed to high prices as power suppliers gamed the system. Only after repeated failures have the designs evolved to something more reasonable, but still far from optimal. The most careful statistical evidence of the importance of design comes not from auction markets *per se* but from the closely related "matching" markets, such as the ones by which most US medical doctors are matched to hospital residency programs. Roth (1991) provides evidence that a certain technical characteristic of the matching rules—whether the outcome lies in the "core"—is a primary determinant of whether certain organized markets succeed in attracting participants over a long period of years.

Successful auction programs need to be well designed in every important respect, of which auction rules are one. Applying an auction theory perspective can be valuable in many ways. It can enable an auctioneer to avoid mistakes like those that marred the 1993 spectrum auction in New Zealand and the 1996 tax auction in Cook County, the flaws of which are obvious in terms of auction theory. It can help the auctioneer to pursue multiple objectives, like promoting minority participation, encouraging alternative suppliers, and enhancing competition among bidders with diverse advantages. Finally, rules can be designed to accommodate complicated preferences and constraints for the bidders and the auctioneer. We will see some examples of that later in this book.

Plan for this Book

This book integrates two projects, which are presented in the next two sections. The first section gives an integrated review of traditional auction theory and is based on courses that I have given over a period of years at Stanford, Jerusalem, Harvard, and MIT. Traditional auction theory is based largely on the theory of mechanism design and the chapter organization follows certain principles of that theory. Much of the analysis is focused on auctions in which each buyer wants only a single object—a condition called "*singleton demand*."

The second section of the book differs from the first in its questions and methods. The questions mainly concern the design of auctions for environments in which there are multiple heterogeneous goods. These environments are fundamentally more complex than ones with singleton demand. One reason is that the number of possible allocations is exponentially larger, which leads to serious issues about the practical feasibility of auction algorithms and bidder strategies. A second is that the case of singleton demand eliminates much of the tension between promoting efficient allocations and ensuring competitive revenues for the seller. In the general case of section II, that tension can be severe. A third difference concerns the problem of *value discovery*. With singleton demand, bidders have only one allocation to evaluate, but in the general case the exponentially larger number of allocations can force a bidder to limit its valuation activities, which can limit both efficiency and price competition.

Because the Vickrey mechanism plays a significant role in both parts of the theory, the next chapter deals with that mechanism.

Auction theory has grown into a huge area of research, and this book reports on only those parts of the theory research that I believe promise to be helpful to auction designers. With that in mind, I have omitted the elegant formal treatments of how auctions perform when there are very many bidders, because these are not directly helpful for choosing among alternative mechanisms.[20] Also omitted are the currently fashionable mechanism design models in which bidders draw extremely subtle inferences from the bids made by others, because these appear to me to lack the robustness needed for a practical design. Finally, I have omitted newer work about subjects like collusion in auctions, revenue-maximizing auctions when goods may be resold, and information processing during auctions. While these are interesting for some applications, knowledge in these areas is evolving too rapidly to be integrated with other parts of the theory.

[20] This research begins with Wilson (1978), includes Milgrom (1979), and especially the beautiful results by Pesendorfer and Swinkels (1997, 2000) and Swinkels (2001).