CHAPTER ONE

Getting to Work

The era of putting auction theory to work began in 1993–1994, 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 1960s, early research had little influence on practice. Since 1994, auction theorists have designed spectrum sales for countries on six continents, electric power auctions in the United States and Europe, CO_2 abatement auctions, timber auctions, and various asset auctions. By 1996, auction theory had become so influential that its founder, William Vickrey, was awarded a Nobel Prize in economic science. In 2000, the US National Science Foundation's fiftieth 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 debate erupted between researchers studying bargaining, who were skeptical that game theory could explain much about bargaining or be useful for improving bargaining protocols, and those investigating in auctions and industrial organization, who believed that game theory was illuminating their studies. Although game theory gained increasing prominence throughout the 1980s and had begun to influence the leading graduate textbooks by the early 1990s, there was no consensus about its relevance 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, which are hard to reconcile with reality. Yet, based on both field data and laboratory data, the contributions of auction theory are hard to dispute. The qualitative predictions of auction theory have been strikingly successful in explaining patterns of bidding for oil and gas[1] and have fared well in other empirical studies as well. Economic laboratory tests of auction theory have uncovered many violations of the most detailed theories, but several key tendencies predicted by the theory find significant experimental support.[2] Taken as a whole, these findings indicate that although existing theories need refinement, they capture important features of actual bidding. For real-world auction designers, the lesson is that theory can be helpful, but it needs to be supplemented by experiments to test the applicability of key propositions and by practical judgments, seasoned by experience.

Whatever the doubts in the academy about the imperfections of game theory, the dramatic case histories of the new auctions seized public attention. An article in 1995 in the New York Times hailed one of the first US spectrum auctions[3] as "The Greatest Auction Ever."[4] The British spectrum auction of 2000, which raised about $34 billion, earned one of its academic designers[5] a commendation from the Queen and the title "Commander of the British Empire." In the same period, game theorists were plying their trade on another important application as well. The National Resident Matching Program, by which 20,000 US physicians are matched annually to hospital residency programs, implemented a new design in 1998 with the help of the economist–game theorist Alvin Roth. By the mid-nineties, thirty-five years of theoretical economic research about fine details of market design was suddenly bearing very practical fruits.

1 See Hendricks, Porter, and Wilson (1994).
2 See Kagel (1995).
3 The design was based on suggestions by Preston McAfee, Paul Milgrom, and Robert Wilson.
4 William Safire, "The Greatest Auction Ever," *New York Times*, March 16, 1995, page A17, commenting on FCC auction #4.
5 The principal designers were Professors Ken Binmore and Paul Klemperer. They give their account of the auction in Binmore and Klemperer (2002). It was Binmore whom the Queen of England honored with a title.

1.1 Politics Sets the Stage

To most telecommunications industry commentators, the main significance of the US spectrum auctions was that a market mechanism was used at all. Spectrum rights (licenses) in the United States 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.[6] 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 sped 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 prizes worth many millions of dollars. Lottery winners were often simple speculators with no experience in the telephone industry and no intention of operating a telephone business. Economic resources were wasted on a grand scale, both in processing hundreds of thousands of applications and in the consequent need for real wireless operators to negotiate and buy licenses from these speculators. 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 in 1993, Congress authorized auctions as the answer. The question of how an auction market for radio spectrum should be designed was left to the Federal Communications Commission (FCC).

1.2 Designing for Multiple Goals

Congress did provide some instructions to the FCC governing the new spectrum auctions. One was that the first auctions were to be begun by July 1994. A second called for the auctions to promote wide participation in the new industry. The FCC initially responded to the second

[6] The process was once characterized by an FCC Commissioner as “the FCC’s equivalent of the Medieval trial by ordeal” (as quoted by Kwerel and Felker (1985)).

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. The statute also specified that the auction process should promote "efficient and intensive use" of the radio spectrum, in contrast with the fragmented use promoted by the lottery system. The meaning of the word "efficient" was initially subject to debate, but it was eventually 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."[7]

There is a powerful tradition in economics claiming that individuals and firms, left to their own devices and operating in a sound legal framework, tend to implement efficient allocations. The argument is that when resources are allocated inefficiently, it is possible for the parties to get together to make everyone better off. So, following their mutual interests, the parties will tend to eliminate inefficiencies whenever they can. This traditional argument has its greatest force when the parties can all see what is required and have no trouble negotiating how to divide the gains created by the agreement. For radio spectrum, with thousands of licenses and hundreds of participants involved, computing just one efficient allocation can be an inhumanly hard problem, and getting participants to reveal the information about their values necessary to do that computation is probably impossible. Compared to the development of a universal standard (GSM) for mobile telephones in Europe, the more fragmented system that emerged in the United States highlights that the lottery system did not lead to efficient spectrum allocations. With so many parties and interests involved, the market took many years to recover from the initial fragmentation of spectrum ownership. During those years, investments were delayed and consumer services degraded. Getting the allocation right the first time does matter. Achieving that with an auction system called for a different and innovative approach.

The FCC, which the law had charged with designing and running the spectrum auctions, had no previous auction experience. Within the FCC, the 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.[8]

[7] Quoted from Vice President Gore's speech at the beginning of FCC auction #4.
[8] Kwerel's initial advocacy is explained in Kwerel and Felker (1985).

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 a wide range of alternatives.

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, John McMillan, to help them evaluate the proposed designs. In the end, Kwerel favored a kind of simultaneous ascending auction, based in large part on a proposal by Robert Wilson and me and a similar proposal by Preston McAfee. The Milgrom–Wilson–McAfee rules called for a simultaneous multiple round ascending auction.[9] This is an auction for multiple items in which bidding occurs in a series of rounds. In each round, bidders make sealed bids for as many spectrum licenses as they wish to buy. At the end of each round the *standing high bid* for each license is posted along with the minimum bids for the next round, which are computed by adding a pre-determined bid increment, such as 5% or 10%, to the standing high bids. These standing high bids remain in place until superseded or withdrawn.[10] An *activity rule* limited a bidder's ability to increase its activity late in the auction, thus providing an incentive to bid actively early in the auction. For example, a bidder that

[9] The principal difference was that the Milgrom–Wilson design proposed the now standard features that bidding on all licenses would remain open until the end of the auction, with progress ensured by Milgrom's activity rule. McAfee's design had no activity rule, and ensured the progress of the auction by closing bidding on each license separately after a period with no new bids on that license.

[10] A bidder that withdraws its bid pays a penalty equal to the difference, if positive, between the eventual sale price for the license and the amount of its withdrawn bid. If the eventual price exceeds its bid, then no penalty is payable.

has been actively bidding for ten licenses may not, late in the auction, begin bidding for eleven licenses.

The theory of simultaneous ascending auctions is best developed for the case when the licenses being sold are substitutes. During the course of the auction, as prices rise, bidders who are outbid can switch their demands to bid for cheaper licenses, allowing effective arbitrage among substitute licenses. One of the clearest empirical characteristics of these auctions is that licenses that are close substitutes sell for prices that are also close – a property that is not shared by most older auction designs.

The initial reception to Kwerel's recommendation was skeptical. The proposed auction was unexpectedly complicated, and FCC Chairman Reed Hundt sought the advice of other FCC staff. He asked the economics staff: If you could pick any design you want, would this be it? He asked those who would have to run it: Can this really work? Even in the short time available to set it up? With the endorsement of his staff, Chairman Hundt decided to take the risk of adopting a new auction design.

1.2.1 Substitutes and Complements

Auctions are processes for allocating goods among bidders, so the challenge of auction design can only be understood by studying the demands of the participants. In the initial PCS auction, there were three groups 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 its own long distance service to offer an attractive and profitable package to consumers.

A second group comprised the existing wireless companies, including AT&T, some regional Bell operating companies, and others. 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 to expand to new regions or perhaps 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 enough of the spectrum to manipulate market prices. To avoid this outcome, the FCC imposed limits on the amount of spectrum that any company could control in any geographic area. These existing wireless operators would be ineligible to bid for a nationwide PCS license of the sort that had typically been awarded in European countries. From MCI's perspective, this meant that a nationwide license might be bought cheaply at auction, so it lobbied the FCC to structure the new licenses in this 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 constrained the possible spectrum allocations. For example, suppose three separate licenses covering areas A, B, and C were put up for sale. If one bidder wanted a license covering A and half of B while the other wanted a license covering C and the other half of B, the license specifications would prevent each bidder from acquiring its optimal allocation. One task of the auction designer was to promote the best (most "efficient") possible allocation, subject to such 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, but be 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 do not 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. For example, the licensee identities 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 much more per license for a package of, say, five or six licenses than for smaller or larger packages.[11] Until such a 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 if 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 differ 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 when licenses might either be substitutes or complements. The impaired performance may take various forms including a loss of efficiency of the outcomes, uncompetitively low revenues to the seller, vulnerability to collusion, complexity for the bidders, and long times to completion.

[11] An instance of this sort arose in the Netherlands spectrum auction in 1998, where 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.

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

1.2.2 New Zealand's Rights 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, 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 consulting firm – NERA – the New Zealand government adopted a *second-price sealed-bid auction* for its first four auction sales. As originally described by Vickrey (1961), the rules of the second-price auction are these: Each bidder submits a sealed bid. Then, the license is awarded to the highest bidder for a price equal to the *second* highest bid, or the reservation price if only one qualifying bid is made. The auction gets its name from the fact that the second highest bid determines the price.

The idea of a second-price sealed-bid 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[12] similar to the one used at Amazon.[13]

In an ascending auction, if a bidder has a firm opinion about what the item is worth, then he can plan in advance how high to bid – an amount that we may call the bidder's *reservation value*. At sites like eBay and

[12] The most common form of an ascending (English) auction is one in which the auctioneer cries out increasing bids and the bidders drop out when they are no longer willing to pay above the current price. The auction ends when there is just one remaining bidder. As the winning bidder is required to pay the current high price, it is optimal for each bidder to stay in the auction only until the current price is equal to his valuation ("reservation value") of the item and not thereafter.

[13] eBay also runs a similar auction, but its fixed ending time involves additional gaming issues as described by Roth and Ockenfels (2000).

Amazon, the bidder can instruct a *proxy bidder* to carry out a *reservation value strategy.* The proxy keeps beating the current highest bid on the bidder's behalf so long as that bid is less than the specified reservation value. 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). If everyone adopts such a reservation value strategy, then the ascending auction is almost the same as a second-price auction.

Strategic considerations in a second-price auction are easy: each bidder should set his reservation value to what the object is worth to him. If it happens that the highest bid among the other bidders is greater than this value, then he cannot do better than to bid his reservation value, because there is no bid he could make that would win the auction profitably. If, instead, it happens that the highest competing bid is less than his value, then setting his reservation value in this way wins and fixes the price at what the competitor bid, which is the best outcome that any bid could achieve. Thus, regardless of the bids made by others, setting a reservation value equal to the bidder's actual value always earns at least as much as any other bid.

The second-price sealed-bid auction has two advantages over most other designs. First, it duplicates the outcome of an ascending bid auction with small bid increments, but without requiring the bidders to be assembled together or to hire agents to represent them in their absence. Second, it presents each bidder with a simple strategic bidding problem: each merely has to determine his reservation value and bid it. This also means that there is no need for any bidder to make estimates of the number of other bidders or their values, for 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 a particular ascending auction. For example, suppose there is such an auction rule with seven identical items for sale, to be awarded to the seven highest bidders in an ascending outcry auction. Again, bidders might sensibly adopt reservation value strategies, bidding just enough to be among the top seven bidders and dropping out when the required bid finally exceeds the bidder's value. An analysis much like the preceding one then 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 value. To duplicate that with a sealed-bid 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.[14]

In New Zealand, the government was selling essentially identical licenses to deliver television signals. 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 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. With licenses that are substitutes or complements, independent auctions inevitably involve guesswork by the bidders that interferes with an efficient allocation.

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. The 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.

[14] The Treasury rule sets a uniform price equal to the lowest accepted bid.

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).

A second problem was even more embarrassing to New Zealand's government officials.[15] 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-bid format to a more standard *first-price* sealed-bid 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 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 risks of acquiring too many licenses and of acquiring too few, 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.

[15] For a detailed account, see Mueller (1993).

1.2.3 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. No bidder would be forced to guess about which licenses to bid on with such an auction. 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 ponder over 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.

1.2.4 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. Because the two licenses had nearly identical technical characteristics and because, for antitrust reasons, no bidder would be allowed to acquire both, 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.[16]

As in the New Zealand case, the main design issue was to minimize guesswork, allowing bidders to choose among substitute licenses based

[16] Dr. Mark Bykowsky of the National Telecommunications and Information Administration (NTIA) was a forceful advocate of the view 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. Nonetheless, 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 for it to be evaluated and adopted immediately.

on their relative prices. When substitute goods are sold in sequence, either by sealed bids or in an ascending auction, a firm 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, during which bidders could place bids on any of the licenses and track bids on all the licenses. The openness of the process would eliminate the guesswork, allowing bidders to switch among substitute licenses, and promote equal prices for licenses that are perfect substitutes.

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, but that raised a new issue. 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 *activity rule*. The general application of an activity rule involves two key concepts: eligibility and activity. A bidder's activity in any round is the *quantity* of licenses on which it has either placed new bids in the round or had the high bid at the beginning of the round. In the example cited earlier, the quantity is just the number of licenses on which a bid is placed, but other quantity measures, including the total bandwidth of the licenses bid or the bandwidth multiplied by the population covered, have also been used. The rule specifies that a bidder's total activity in a round can never exceed its eligibility. A bidder's initial eligibility, applicable to the first round of the auction, is established by filing an application and paying a deposit before the bidding begins. Its eligibility in each later round depends on its recent bidding activity. One simple form of the rule specifies that a bidder's eligibility in any round after the first is equal to its activity in the preceding round. Thus, bidders that are not active early in the auction lose eligibility to place bids later in the auction. This rule speeds the auction and helps bidders to make reliable inferences about the remaining demand at the current prices.

The FCC rules have evolved since the original 1994 design, but larger changes have been made to adapt the simultaneous ascending auction to other applications. One common variation arises when there are many units of each kind of item, such as auctions involving the sale of electricity contracts. In these auctions, for each item, each bidder bids its quantity demanded at the current price indicated on a *clock* visible to all bidders. The clock starts at a low price and keeps raising the price at any point at which the current total demand of all bidders exceeds the supply of that item. When demand equals supply on all items, the auction ends. A series of such clocks record the current prices for the various goods, and the rate of movement in these clocks 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 share several key characteristics with their FCC ancestor. Bidding on all items takes place simultaneously, so bidders can respond to changing relative prices. Prices rise monotonically, ensuring that the auction progresses in an orderly and predictable way. All bids are serious and represent real commitments. There is an activity rule that prevents a buyer from increasing its overall demand on all items 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.

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 at the same time – January 2002 for the first sale. Because different buyers wanted different mixes of contract lengths and because all contracts covered the first quarter of 2002, EDF regarded the different kinds of contracts as substitutes.

Lawrence Ausubel and Peter Cramton developed the auction design. The first step was to assist EDF in developing a standard for "scoring" bids on contracts of different lengths. Bids expressed a price per megawatt per month that the buyer would pay for the right to acquire power. For the initial auction, EDF specified that the price for a three month contract for base-load power would always be €2139 higher than the corresponding price for a six month contract. Similarly, price differences were specified between the three-month contract and contracts lasting ten, twelve, twenty-four or thirty-six months. During the auction itself, the price clocks were controlled to maintain these price relationships, for example, the price of a three-month contract was at all times €2139 higher than the price of a six-month contract. Prices for contracts of all lengths continued to rise until the total remaining demand exhausted the total power available for the initial three-month period.[17] Such an auction creates competition among bidders for contracts of different lengths, increasing both efficiency and sales revenue compared to more traditional auction designs. Recently, the EDF auction has been further modified to include a "supply curve," so that total quantity of power sold depends on the price level.

1.3 Comparing Seller Revenues

The question most frequently asked of auction designers is: What kind of auction leads to the highest prices for the seller? The answer, of course, depends on the particular circumstances, but even the thrust of the answer surprises many people: There is no systematic advantage of either sealed bid over open bid auctions, or the reverse.

A particular formal statement of this conclusion is known as the *payoff equivalence theorem*. It holds that for an important class of auctions and environments, the average revenues and the average payoffs of bidders are exactly the same for every auction in the class. 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, then in theory B will bid just enough to ensure that it wins – around $10 or $10.01 – and A will likely bid close to $10. If they behave that way, the price will be just the same as in the ascending auction.

The argument in this simple form was first made by Joseph Bertrand (1883). Nearly a century later, William Vickrey observed, that a similar conclusion holds on average for a much wider class of auction rules and in a more realistic set of situations than the one described here. For

[17] For example, in the sale of power beginning January 2002, when the total demand exceeded the power available for the first quarter of 2002, the auction ended. Any remaining unsold power for, say, the second quarter of 2002 was then included in subsequent sales.

forecasting average revenue, it is irrelevant which auction is used, within a certain class of standard auction designs.

Experienced auctioneers often contest this irrelevance conclusion. Those who advocate ascending auctions argue that they generate more excitement and more competition than sealed bids. After all, they claim, 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 bids counter by arguing that ascending auctions never result in more being paid than is absolutely necessary to win the auction; there is no money "left on the table." Sealed bids frequently result in lots of money left on the table. For instance, 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 bid was about 60% higher than the second highest bid, so 40%, or about \$1 billion, was left on the table.[18]

Similar arguments among practitioners arise quite frequently, sometimes with variations. In the United States, the staff of the Treasury Department have periodically argued the relative merits of two alternative auction schemes for selling bills. In one scheme, each bidder pays the amount of its own bid for each bill it buys; in the other, all bidders pay the same *market-clearing price,* identified by the lowest accepted bid. Advocates of the first ("each pays its own bid") scheme say that the government will get more money from the auction, because winning bidders are by definition people who have bid more than the lowest acceptable bid. Advocates of the second ("uniform price") scheme counter that bidders who know they must pay their own bid when they win will naturally bid less, reducing the market-clearing price and leading to lower revenues.

Informal arguments like these show that the matter is subtle, but they do not settle the issue. A formal analysis based on the *payoff equivalence theorem* discussed in chapter 3 helps to cut through the confusion. Under certain idealized conditions, if the allocation of lots among bidders is the same for two different designs, then the average payoffs to all parties, including the average prices obtained by the seller, must also be exactly

[18] Although 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.

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.

The practical uses of the payoff equivalence theorem are similar in kind to the uses of the Modigliani–Miller theorems in financial economics, the Coase theorem in contract theory, and the monetary neutrality theorems in macroeconomics. All of these theorems assert that under idealized conditions, particular effects cannot follow from identified causes.[19] For example, according to the Modigliani–Miller theorems, if decisions about debt–equity ratios and dividend policies merely slice up the total returns to a firm's owners without affecting the firm's operations, then those decisions cannot affect the firm's total market value. Today, financial economists explain financial decisions by focusing on how financial decisions might affect a firm's operations – its taxes, bankruptcy costs, and managerial incentives. Similarly, according to the Coase theorem, if there were no costs or barriers to transacting, then the default ownership of an asset established by the legal system could not affect value. Today, economic theorists explain features of organization in terms of costs and barriers to transacting, including incomplete information and incomplete contracts. The payoff equivalence theorem is similar: the payment terms of an auction do not affect the seller's total revenue unless they are associated with a change in the allocation of the goods. Today, analysts focus more attention on how assumptions of the theorem are violated and the consequences of those violations or, for government regulators, about the implied trade-offs between their allocation and revenue objectives.

The planning for a sale of electrical power in Texas in 2002 illustrates how the payoff equivalence theorem has been applied in practice. According to the planned auction design, the auctioneer would gradually raise the prices for any products with excess demand and would accept quantity demands from the bidders, in much the fashion that Leon Walras once described. The auctioneer would not tell the bidders the

[19] According to the Modigliani–Miller theorems, under its idealized frictionless-markets conditions, a firm's financial structure and dividend policy cannot affect its market value. According to the Coase theorem, under other idealized conditions, the initial allocation of ownership rights cannot alter the efficiency of the final allocation. Monetary neutrality theorems hold that under yet other idealized conditions, monetary policy cannot change real outcomes in an economy. The payoff equivalence theorem holds that under its idealized conditions, changing payment rules cannot affect the participants' final payoffs.

quantities demanded by others. The rules called for the auctioneer to stop raising the price for a product when its total demand falls to the level of available supply. Texas ratepayers benefit from the revenues of this power sale, and the ratepayers' advocate argued that the auctioneer should continue to raise prices until demand is actually *less than* supply, and should then roll back the price by one increment. The idea was to sell power for the *highest* market clearing price, rather than the lowest one. This rule was problematic for a variety of reasons relating to the details of the auction, and the design team cited the payoff equivalence theorem to argue that there was little reason to expect that the proposed change would lead to higher prices on average, because bidders would bid differently if the payment rules were changed. A bidder that knows it may acquire power at a lower price if it withdraws demand early will be more inclined to do so than a bidder that knows it cannot cause a price rollback. The net effect on revenues is hard to predict, because it depends on how the proposed new rule changes the allocation. Eventually, the ratepayer advocate agreed not to oppose the auction design.

1.4 The Academic Critics

Economists who are 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 generate ideas using theory, test those ideas when they can, and implement them with awareness of their limitations, supplementing the economic analysis with worst case analyses and other similar exercises.

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.

1.4.1 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 does not matter at all. After all, say the critics, once the licenses are