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, power auctions in the US and Europe, and several smaller 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 been used for worldwide sales totaling more than $100 billion.

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 appeared 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 those fields. Although game theory had risen in status 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 of game theoretic predictions about bidding behavior, but key tendencies predicted by the theory do find experimental support. The findings suggest 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 attracted considerable public attention. A New York Times article hailed one of the early US spectrum auctions as "The Greatest Auction Ever,"[1] and one academic auction designer—Professor Ken Binmore—was named a "Commander of the British Empire" for his work: the British spectrum auction conducted in 2000 using a variant on the earlier US design[2] raised more money than any other auction in history. By the mid-nineties, thirty years' worth of economic auction research had arrived all at once on a bustling market scene.

---

[1] William Safire, "The Greatest Auction Ever," *New York Times*, March 16, 1995, page A17.

[2] Designed by economists Ken Binmore and Paul Klemperer.

Politics Sets the Stage

Not all of the public emphasis was placed on the auction designs themselves. To some telecom industry commentators, the main significance was that a market mechanism was used at all. Spectrum rights (licenses) in the US had long been assigned using *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.[3] 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.

Not everyone favored the change. As auction prices of spectrum soared, large European telephone companies (“telecoms”) began to conjure arguments against the auction processes, claiming that auctions harm consumers and damage economic welfare. Two fallacious arguments were commonly raised. According to one, companies that pay high prices to acquire spectrum rights will wind up recovering those through higher consumer prices of wireless services, or will be willing to spend less on advanced services for consumers. Readers schooled in economics will recognize this argument as a “sunk cost fallacy”: once the licenses have been paid for, the price paid for the spectrum has no bearing on what the telecom can most profitably charge or what services it can most profitably offer. In the US, latecomers to the wireless communications market paid the highest prices for their spectrum licenses and were then most aggressive in reducing prices and introducing new services to attract their initial customers.

The second argument made by opponents of auctions is that if companies pay actual market prices for spectrum, they will have little left over to spend on building the physical system, delaying the delivery of valuable services to consumers. The very high

---

[3] 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)).

prices bid for spectrum in the British auctions and the subsequent financial difficulties faced by winning bidders are offered as evidence.

This amounts to a claim that the government should offer huge subsidies to telecoms to help them build their expensive wireless systems quickly, rather than spending the same public moneys on welfare, health, education, defense or tax cuts. In a mostly capitalist economy with sound financial markets, it would be inconsistent to subsidize large, publicly traded companies so that they can more quickly supply a profitable, private service. The experience of the winning bidders in the British auctions simply shows that they paid more for the licenses than they were actually worth. Overspending is hardly a convincing reason to funnel public money to a company.

Designing for Efficiency

When the US Congress authorized the first spectrum auctions as part of the 1993 omnibus budget bill, it included several explicit instructions. One was that the first auctions were to be run in that fiscal year. A second specified that the purpose of the auctions was to achieve "efficient and intensive use" of the radio spectrum. 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."[4]

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.[5]

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

---

[4] Quoted from Vice President Gore's speech at the beginning of FCC auction #4.

[5] Kwerel's initial advocacy can be found in Kwerel and Felker (1985), "Using Auctions to Select FCC Licensees," OPP working paper 16.

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[6] 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.[7]

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.

---

[6] Professor John McMillan.

[7] These rules were based mainly on a detailed proposal that I prepared in collaboration with Professor Robert Wilson, including an *activity rule* of my design to pace the auction and make it possible to prescribe that bidding would remain open on all licenses until the final round. The FCC rules also resembled an alternative simultaneous ascending auction proposed by Professor Preston McAfee, which differed mainly in its license-by-license ending rule.

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.[8] *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*.[9] Economic discussions of the auction design are usually organized by emphasizing these two 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

---

[8] 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.

[9] 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.

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 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

---

[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.

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, on the advice of its consultants, did not adopt this “highest rejected bid” rule. Instead, it chose 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. When the values are so interdependent, unlinked, independent auctions inevitably involve guesswork that gets in the way of an efficient allocation.

---

[11] The Treasury rule set 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).

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 five 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

---

[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 was too risky to run immediately.

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 prices increase. This rule ensures plenty of activity early in the auction and enhances information exchange: bidders can tell, by declining bidding activity, when the auction must be nearing its end.

The original FCC design has been modified and adapted in various ways for other applications. A 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. In this design, as in all the FCC progeny, there is an activity rule that prevents a buyer from increasing its overall demand as prices increase.

New variations based on the original design continue to be created to solve a wide range of economic problems. One particularly interesting variation was used recently by Electricité de France[13] in a pair of sales of electrical power contracts. This particular application allowed substitution not only by the buyers of power but also by the seller. The sales, which took place in September and December of 2001, were for based-load and peak-load power contracts for various terms (3-month, 6-month, etc), all beginning in January 2002. From the seller's perspective, the total of all its contract obligations beginning at that date, regardless of duration, was limited by the quantity it could sell in the first quarter, so the contracts were substitutes. During the auction, the auctioneer raised the price clocks for the power contracts of various durations in a pre-determined way, based on the recent relationship between contracts of different lengths. It stopped raising prices when the total remaining demand was equal to total supply. Such an auction creates competition among bidders for contracts of different lengths, adding to the efficiency of the outcome.

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 *revenue 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

---

[13] Larry Ausubel and Peter Cramton, professors at the University of Maryland and officers of Market Design, Inc, created this design.

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.[14]

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

---

[14] 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.

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.[15]

Important Criticisms

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.

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.[16] Auction and bargaining theory and the history of

---

[15] 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.

[16] 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—and the assumption that the

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.[17]

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 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

---

parties have unlimited budgets, so spending on spectrum rights does not impair the ability to invest in infrastructure. We omit these factors here because they appear to be of secondary importance for resolving the question of how auction design can matter.

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

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. The mechanical engineer pays attention to factors that increase friction and builds in redundancy and safety margins, and the 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, the issues that most auction design models capture form only a small subset of the issues that a real auctioneer may face. 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, the sum of these considerations prevent the proper auction design from being completely determined by solving an optimization problem.

- *What?* 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.
- *Fighting collusion?* The European spectrum auctions, with their very high stakes, provided some interesting examples of before-the-auction actions to reduce competition. In Switzerland, last minute changes resulted in only four bidders showing up for four spectrum licenses. The British auction designers favored leaving the auctioneer leeway to ensure that the number of licenses would always be less than the number of bidders. In Italy, reserve prices were set very high to limit the benefits of potential collusion.

---

[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.

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

The mechanism design purist's view, which holds that the only consistent approach is to develop theoretically "optimal" mechanisms, is much too extreme to be useful in practice. Our models of human behavior are not nearly accurate enough for 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. Business procurement officers ask whether a situation is "*auctionable*," when the alternatives are to bargain individually with suppliers or simply to accept their posted prices.

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. Successful experiments testing the theorist's proposals[19] played a critical role in convincing the FCC to adopt the theoretically motivated design, although other factors were important as well.[20] 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

---

[19] Charles Plott designed and conducted these experiments in his lab at the California Institute of Technology.

[20] Implementation issues also played a huge role in the debate. The very possibility of running a computer implemented simultaneous auction drew heckles 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.

indispensable. Theory guides the design of experiments, suggests which parts of any 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 Issues

The final criticism is that, in the real world, the most important issues are ones about which economic theory has very little to say. These are issues of implementation.

One example of such an issue is the bidder interface. In the original FCC auction software, bidders would enter bid amounts directly into a computer interface. There were several occasions on which bidders would make what came to be called “fat finger bids.” When trying to bid $1,000,000, they might accidentally bid $10,000,000—an error encouraged by the fact that the early interfaces could not accept commas in the bid field.

A related problem with the bidding field is that bidders entering very large numbers would sometimes use the last digits to communicate messages to other bidders. For example, if bidder A wishes to retaliate against bidder B for bidding on license 468, it might raise the price of another license on which B has the current high bid of, say, $9,000,000 by bidding $10,000,468.

In practice, the FCC solved these two problems by changing to an interface in which bids must be selected from a drop-down menu. That eliminates badly mistaken bids and greatly reduces any bidder’s ability to send messages with its bids.

The scope of these practical issues is quite extensive. Besides bidder interfaces, there are issues of defining the packages or “lots” to be sold, determining the bid increments, specifying how much time bidders have to bid and respond, whether the auction will close at a fixed time or whether the time will depend on bidding activity, and so on.

Although economic theory is silent on some of these matters, we shall see that the theory does lend insights into several of the practical questions. The closing time, the bid increments, and the definitions of lots are all subjects that have been usefully analyzed.

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 the 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 both in its main questions and in its primary methods. The questions concern the design of resource allocation auctions for environments in which there are multiple heterogeneous goods. These environments are fundamentally more complex than ones with singleton demand for several reasons. First, the number of possible allocations is exponentially larger, which leads to serious issues about the practical feasibility of auction algorithms and bidder strategies. Second, in the case of singleton demand all goods are substitutes, which eliminates 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. Third, there are problems 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.

Some of the same issues arise in section I in connection with the Vickrey mechanism, but the analyses of section II emphasize a different approach. The analytical style is drawn mainly from matching theory, and the section is named after that approach.

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

Section I:
The Mechanism Design Approach

The four chapters of section I apply "mechanism design theory" to problems of auction design. The main concepts of mechanism design theory can be described succinctly in intuitive language, which provide us with a convenient point of departure. These intuitive descriptions conceal technical details that are occasionally important, so the mathematical development is indispensable for a full understanding of the theory.

Mechanism design theory distinguishes sharply between the apparatus under the control of the designer, which we call a "*mechanism*," and the things that are beyond the designer's control, which we call the "*environment*." A mechanism consists of any rules that govern what the participants are permitted to do and how these permitted actions determine *outcomes*. An environment comprises three lists: a list of the participants or potential participants, another of the possible outcomes, and another of the participants' possible *types*—that is, their capabilities, preferences, information and beliefs.

At this level of abstraction, the theory can be applied in many different kinds of environments. In a model of political mechanisms, the participants might consist, for example, of legislators, regulators, or the electorate. An outcome could be a set of bills that are enacted, regulations that are adopted, or officials that are elected. The mechanism analyst might be concerned about the electoral system affects the likelihood of legislative stalemate, the power of certain minority interests, or how the electoral system distorts choices by politicians concerned with reelection. A typical problem might be to find the system that makes it most likely that the mechanism implements, say, the optimal utilitarian policy choice.

In economic mechanism models, participants might be workers, firms, shareholders, managers, or the husband and wife in a family. Outcomes might describe resource allocations, job assignments, the sharing of housework or allocation of the family budget, etc. The most commonly studied mechanisms in economic theory are *resource allocation mechanisms* in which an outcome is a resource allocation, that is, a specification of the bundles of resources assigned to each participant and of the levels of certain public goods, like parks, libraries, or police and fire services.

Mechanism design theory has historical roots in economics as part of the theory of how a benevolent government should run its affairs. A benevolent government is one that wants to implement a "good" outcome, but faces certain obstacles. For example, it may lack information about people's resources, capabilities, and preferences, so that it does not know what outcomes are feasible, how to implement them, or how citizens rank the possible outcomes. Even if it knew those things, it might be unable to tell whether people are following its instructions. These are *information problems*, and solving such problems optimally is a traditional focus of the theory.

The mechanism designer may also face a *commitment problem*. For example, if the designer is someone who runs a firm then, despite any promises the designer may make, a worker fear she will change a quota adversely if he—the worker—reveals that he is capable of increasing production substantially.

Both kinds of problems play a role in mechanism design theory and in its application to the *economic theory of contracts*. For auction theory, however, the main focus of our attention will be information problems, in which we assume that the bidders know more about their values than whoever is designing the auction rules.

A central assumption of the mechanism design approach is that nobody cares about the process itself, except insofar as it determines the outcome. In the real world of auction design, one reason that analytically optimal designs may be rejected is that they seem too complicated or unfamiliar. Despite the loss that comes from assuming away preferences about processes, this assumption provides a solid and tractable structure on which to begin to build a formal analysis of mechanisms.

Mechanism theory evaluates alternative designs based on their comparative *performance*. Formally, performance will refer below to the function that maps environments into outcomes. "When it rains, we distribute umbrellas; when the sun shines, we distribute bathing suits" is better performance than the opposite distribution pattern. This illustrates a performance function, because it is a mapping from environments to outcomes.

The goal of mechanism design analysis is determine what performance is possible and how mechanisms can best be designed to achieve the designer's goals. The three most common questions are: Is it possible to achieve a certain kind of performance, for example to identify a mechanism that leads to efficient allocations of goods in *every* environment in some class? Precisely which performance functions are *implementable* by some mechanism? What mechanism *optimizes* the performance (according to whatever performance criterion that the mechanism designer may have in mind)?

An *auction* is a mechanism to allocate a set of resources to maximize the total cash amounts offered by buyer-bidders (or to minimize the total cash amount demanded by seller-bidders). Auction models include a description of the potential bidders, the possible resource allocations (describing the number of goods of each type, whether the goods are divisible, and whether there are legal or other restrictions on how the goods may be allocated), and the values of various packages of goods to each participant. Values may be determined in subtle ways. For example, when a bottle of fine wine is sold at auction, the winning bidder's payoff may depend on how much she likes the particular wine, or likes the prestige of winning the bottle, or likes keeping the bottle away from a certain competing collector. Losers, too, may care about the outcome, for example because they expect that if a certain friend wins the bottle, she will serve it at an upcoming wine tasting party. The mechanism designer's problem is to determine the rules of the auction—what bids are allowed, how the resources are allocated, and how prices are determined—in a way that achieves some objective, such as maximizing the seller's proceeds.

Mechanism design theory was initiated by three applications to important problems in economics. The first was William Vickrey's design of auctions that achieve efficient allocations of resources in a wide range of circumstances. That theory is reviewed in detail in the next chapter. The second was the Vickrey-Mirrlees design of an optimal income tax and welfare system given a utilitarian objective. Vickrey formulated the problem to incorporate the ideas that individual utility depends on income and leisure, that different people have different opportunities to generate income by sacrificing leisure, that the taxing authority can observe only total income, and that the tax system affects individual willingness to sacrifice leisure to generate income. The problem was to create an income tax and welfare system in which the government transfers income from some people to others to maximize a utilitarian objective: the total utility of everyone in society. The solution would presumably entail transfers from those with high earning ability to those who are relatively disabled, that is, who have low earning ability.

The optimization problem implied by this formulation was later revisited and solved by James Mirrlees, whose methods of analysis and auxiliary assumptions were mimicked by subsequent researchers. For their efforts on these two problems, Vickrey and Mirrlees shared the 1996 Nobel Prize in economic science.

The third application was the Clarke-Groves analysis of the optimal provision of public goods, in the economy or with other organizations. The methods and conclusions of the analysis are quite similar to Vickrey's auction theory; consequently it is most convenient to treat the two together in the next section.

In the years that followed, the same techniques were applied to public sector problems, like the optimal state regulation of public utilities to maximize consumer welfare, and to private sector problems, like the optimal design of contracts to maximize the welfare of the contracting parties. Roger Myerson initiated the application of these ideas in modern auction theory, in a study of how auctions could be designed to maximize the seller's expected revenue.

Formalities of the Mechanism Design Model

Mathematically, an *environment* is defined to be a triple $(\mathbf{N},\Omega,\Theta)$, whose elements are $\mathbf{N} = \{1,\dots,N\}$—the set of *participants* (or potential participants) in the mechanism, Ω—the set of possible outcomes over which the participants and the mechanism designer have preferences, and $\Theta = \Theta^1 \times \dots \times \Theta^N$ —the set of *type profiles* $t = (t^1,\dots,t^N)$, where a participant i's *type* (t^i) characterizes her information, beliefs, and preferences. Often, we denote type profiles by (t^i,t^{-i}), where t^{-i} is the type profile for the participants besides i. A (*strategic form*) *mechanism* is a pair $(\mathbf{S},\omega)$ where $\mathbf{S} = \mathbf{S}^1 \times \dots \times \mathbf{S}^N$ is the set of possible *strategy profiles* ($\mathbf{S}^j$ is the set of possible strategies of a typical player j) and $\omega : \mathbf{S} \to \Omega$ maps strategy profiles to outcomes.[21]

The content of type profiles also varies among applications. At a minimum, types describe individual preferences over outcomes. Usually, this is modeled by treating the type as a parameter of some utility functions, writing $u^i : \Omega \times \Theta \to \mathbb{R}$ to describe these utilities. This notation allows the possibility that individual i's utility may depend on all the participants' types. For example, in Akerlof's famous "lemons" model of the market for used cars, the buyer's utility from acquiring the car at a given price depends on the

---

[21] This is a *strategic form* description of the mechanism. Mechanisms can alternatively be described in *extensive form*, which entails a complete description of the succession of possible moves (the "*game tree*"), what each player learns before it moves (the "*information sets*"), and the outcome that follows each possible sequence of moves. The difference between the two descriptions is potentially significant when the analyst applies an extensive form solution concept, such as *sequential equilibrium* or *perfect equilibrium*.

car's condition, which only the seller knows. A similar structure arises in every mainstream model of "adverse selection."

With only a few exceptions, the theory of mechanism design has abstracted from the possibility of adverse selection to focus on other matters. This abstraction is accomplished by limiting attention to the *private values* case, in which each participant's utility depends only on its own type: $u^i(x,t)=u^i(x,t^i)$. In this case, a participant's ranking of the outcomes in Ω cannot be affected by anything it may learn from the other participants. Except where specifically noted, all the mechanisms in this chapter deal with the private values case.

Most mechanism models assume that participants are uncertain about what other participants know. In *Bayesian* models, what each participant *i* believes about the others depends on its own information and is described by the conditional probability distribution $\pi^i(t|t^i)$. Throughout most of this chapter, we employ the *Harsanyi doctrine* that the beliefs represent posterior beliefs derived from a *common prior* distribution π. In that case, $\pi^i(t|t^i) = \pi(t|t^i)$ obeys the laws of conditional probability. Although this doctrine is restrictive and rules out certain interesting and realistic phenomena, it is also an effective way to abstract away from "*betting pathologies*."[22] To focus attention on other matters, the Harsanyi doctrine is widely used in mechanism design models.

For each mechanism that may be chosen and each realization *t* of the type vector, a corresponding strategic form *game* can be defined. Formally, the game $\left(\mathbf{N},\mathbf{S},U(\cdot\,|\,t)\right)$ is a triple consisting of a set of players, a set of strategy profiles, and a *payoff function* mapping strategy profiles into payoffs. The payoff function is given by: $U^i(\sigma^1,\ldots,\sigma^N,t)=u^i(\omega(\sigma^1,\ldots,\sigma^N),t)$. In case the players are Bayesians, adding the beliefs as described above completes the description of a Bayesian game.

Given a mechanism $(\mathbf{S},\omega)$, if the game theoretic solution concept forecasts that a particular strategy profile $\sigma=\left(\sigma^1(t^1),\ldots,\sigma^N(t^N)\right)$ will be played, then one can use that forecast to predict and evaluate the performance of the mechanism. The forecasted outcome will be $x(t)=\omega\left(\sigma^1(t^1),\ldots,\sigma^N(t^N)\right)$. The function *x* mapping type profiles to outcomes is the *performance function* corresponding to the mechanism $(\mathbf{S},\omega)$.

---

[22] Legend has it that the "betting pathology" was first discovered in the coffee room of the Stanford University economics department, when Professors Joseph Stiglitz and Robert Wilson disagreed about whether a certain uncomfortable seat cushion was stuffed with foam or feathers. They agreed to bet $10 on the issue and to cut open the cushion, with the loser to pay for a new cushion. Alas, the department administrator stopped them before they could execute their agreement. Notice that this agreement, which each participant regarded as a good deal for himself, required the destruction of real resources to no productive purpose. Any mechanism that does not involve such side bets can always be Pareto-improved by adding a bet that makes each participant strictly better off, while paying the mechanism designer a fee for arranging the bet. Adopting the Harsanyi doctrine abstracts from many such side bets to focus attention on other aspects of the mechanism design problem. As we shall see, however, even with the Harsanyi doctrine, there can still be room for Pareto-improving side bets when the participants' types are statistically correlated.

Many game theoretic solution concepts are not single-valued; for example, many games have multiple Nash equilibria. There are several ways to proceed at this juncture, but for this chapter we focus just on the following one. When a game has multiple solutions, we define the *augmented mechanism* $(\mathbf{S},\omega,\sigma)$ to be the mechanism plus a selected solution. The idea is that the solution σ represents a recommendation made by the mechanism designer to the participants. If the recommendation is consistent with a solution concept that adequately captures the participants' incentives, then no participant would have any incentive to deviate from the recommendation and, it is argued, σ is therefore a reasonable prediction of how the participants will behave.

When σ is a solution according to some solution concept, we say that the mechanism $(\mathbf{S},\omega)$ or the augmented mechanism $(\mathbf{S},\omega,\sigma)$ "*implements*" the performance $x=\omega\circ\sigma$. Sometimes, we attach the name of the solution concept to modify the verb, saying that a mechanism "*implements in dominant strategies*" or "*Bayes-Nash implements*" the particular performance.

The Chapters

We move through the mechanism design approach to auction theory in a series of steps. In chapter 2, we review the Vickrey analysis of auctions and the related Clarke-Groves analysis of public decisions. The Vickrey-Clarke-Groves (VCG) design establishes a useful benchmark to which subsequent analyses of resource allocation mechanisms must be compared.

Chapter 3 extends the analysis by introducing the Milgrom-Segal extension of the classic envelope theorem, which describes payoffs in integral form for easy application to problems in mechanism design.[23] The chapter uses the envelope theorem to prove a *payoff equivalence theorem*, which in turn leads quickly and easily several of the main results of mechanism design theory and auction theory.

Chapter 4 continues the analysis by adding the assumption that individual preferences satisfy one or another "single crossing property." This leads to the *constraint simplification theorem*, which allows both a short and convenient analysis of standard auction designs and an easy characterization of which performance functions are implementable by some mechanism. With additional simplifying assumptions this leads to the celebrated "optimal auction" theory, but it also leads to some interesting comparisons of different auction designs when the simplifying assumptions are relaxed.

---

[23] Most of the literature on mechanism design instead starts with the "*revelation principle*," which asserts that a performance function is implementable (in dominant strategies or Bayes-Nash equilibrium) if and only if it is incentive-compatible and uses it to derive specialized versions of the integral form envelope theorem. The arguments in this book are simplified and the results strengthened by proving instead a general envelope theorem and using the revelation principle in a more limited way.