CHAPTER SEVEN

The Greatest Auction in History

R. Preston McAfee, John McMillan, Simon Wilkie

In August 1993 President Bill Clinton signed a historic law granting the U.S. Federal Communications Commission (FCC) the authority to auction spectrum licenses.[1] The origin of this law dates back to Ronald Coase's 1959 proposal to sell the radio spectrum. Congress gave the FCC until August 1994 to begin the first auction. To someone inexperienced in the activities of large bureaucracies, a year sounds adequate to design and operate an auction, but it is remarkable that the FCC was able to meet this requirement and commence its first auction on July 25, 1994. In order to run an auction, the FCC needed to choose an auction design, and in order to do that, it was required by law to provide adequate time for persons to file formal comments and reply comments on its proposed procedures. These comments, along with staff recommendations, were used by the FCC commissioners to determine an auction form that would survive judicial challenge. Finally, the FCC staff then implemented the commission's decision, relying on outside contractors to develop auction software. Timely implementation was particularly challenging given that the commission decided on a novel auction method, the electronic simultaneous multiple-round bidding auction, based on economic advice.

Congress and many economists underestimated the number of choices in designing an auction. In modern business parlance, auctions have many "dials and levers" that can be used to tweak performance and outcomes, but the presence of these choices creates an enormous economic and political challenge. This chapter describes how that challenge was met using modern game theory. The efforts of economists led to what William Safire (1995) called "the greatest auction in history," using methods that have been copied around the world to sell over U.S. $100 billion in radio spectrum.

Auctions were not seriously considered for awarding spectrum licenses until two other methods ran into difficulty. The bureaucratic mainstay, administrative allocation by comparative hearing, began to break down as the value of the items to be allocated increased. Hearings and court challenges take years and incur large costs for both participants and the government. Congress replaced these cumbersome administrative procedures with lotteries in 1982. The potential for a windfall gain from winning a license resulted in nearly 400,000 applications for cellular licenses. With so many applications, the magnitude of the total private expenditures in obtaining licenses (rent seeking) was large for a process that only by accident assigns licenses to parties that value them most highly.

When Congress mandated the use of auctions, expert academic economists became involved primarily because neither the potential bidders nor the FCC had experience in auctions. The FCC hired John McMillan as a consultant to advise it on auction design; he supplemented FCC economist Evan Kwerel, the driving force behind the auctions. The major bidders hired economists, including Paul Milgrom, Robert Wilson, Preston McAfee, Barry Nalebuff, Peter Cramton, David Salant, and Rob Gertner, to write comments to file with the FCC on auction design and to advise on bidding strategy during the auctions.

The absence of a history of auctions at the FCC was a significant advantage to the commission in developing a novel and appropriate design based on recent advances in auction theory. It is incredibly difficult for a bureaucracy to change an existing methodology that works moderately well, even if an alternative is clearly superior. Change involves career-destroying risks and opens lobbying doors that might result in worse choices. In the design of the Personal Communication Services (PCS) auctions, the absence of an existing default meant that all designs were on the table. Moreover, the use of academic economists helped focus the debate on the public interest rather than on the firms who in most cases paid their bills. Academic economists were generally unwilling to defend their clients' goals when these goals conflicted with the social interest. In addition, the determination of FCC chairman Reed Hundt to "do the right thing" helped prevent the agency capture all too common in Washington, a determination strongly aided by Kwerel's analyses.

FCC Auction Issues

Congress mandated that the FCC auctions should

- promote efficient and intensive use of the electromagnetic spectrum;
- promote rapid deployment of new technologies;
- promote economic opportunity and competition by dissemination of licenses to a wide variety of applicants;
- recover for the public a portion of the value of the public spectrum resource.[2]

Economists emphasized four goals: simplicity, efficiency, revenue, and diversity. Although simplicity was not a goal specified by Congress nor found in economics texts, economists recognized immediately that because the spectrum auctions would be conducted for the first time with inexperienced bidders, the auction design needed to be simple to understand. A complex design was likely to be misunderstood, leading to unintended choices and inefficient outcomes.

A second aspect of simplicity, and one harder to implement, requires that a simple strategy be optimal or nearly optimal behavior in the auction. For example, "bid your estimated value" is a reasonable strategy in a Vickrey auction[3] but a terrible strategy in a first-price (winners pay their bid prices) sealed-bid auction. A first-price sealed-bid auction has the property that bidding sensibly requires assessing the bidding strategies of rivals, a hard thing to do in a static game. Thus the first-price sealed-bid auction, although simple to understand, has no sensible simple strategy for participants. Economists were very much concerned to articulate simple bidding strategies for bidders that would perform well. It was expected that novice bidders would probably adopt such strategies.

Congress mandated some form of efficiency, which economists generally interpreted to mean that the bidders with the highest value should obtain the license. If there had been only one license, efficiency would not have been a significant challenge. However, thousands of licenses were available, spanning different geographic regions and distinct spectrum bands. Efficiency therefore meant both figuring out which aggregations of licenses made sense and doing so without substantial delay, because delay itself is a cost. With the exception of some work on bundling, economic analysis provided little guidance. Efficiency requires appropriately defining licenses both geographically and across spectrum bands and designing the auction mechanism to award those licenses. In defining the geographic scope of licenses, one must consider the ability of the auction mechanism and secondary markets efficiently to aggregate small licenses into larger coverage areas versus the ability of secondary markets to efficiently disaggregate large licenses into smaller areas.

Before the broadband PCS auctions, some economists argued for dividing the country into large regions (or even the entire nation). The problem with this band plan was that most of the potential bidders owned some portion of the existing cellular spectrum and generally would be seeking to buy a complement of their existing holdings. Indeed, limits on ownership precluded them from buying in areas where they already had holdings. Large regions would not permit purchase of complementary holdings in the auction and would lead to few bidders or requiring divestiture of existing holdings by participants.[4]

Other economists argued for dividing the nation into many small geographic areas without proposing an auction that would be likely to facilitate efficient aggregation. These economists relied on secondary markets to achieve efficiency. On the basis of experience with administrative assignment and lotteries, we know that high transaction costs in secondary markets can delay and in some cases prevent altogether the efficient aggregation of spectrum. The cellular spectrum allocated earlier remained a patchwork of ownership, with incomplete roaming agreements, even twelve years later. History suggests that getting the initial allocation substantially right was very important for the speed of comprehensive deployment of wireless services.

A major challenge to efficient allocation is the presence of local synergies in value. Local synergies generally arise from three distinct motivations—creation of seamless roaming, elimination of boundary interference, and advertising or scale economies in management. Ausubel et al. (1997) find evidence of such synergies in the bidding. The history of spectrum license aggregations suggested that it was difficult to forecast the shape of efficient assignments. GTE created an aggregation of licenses in the Southwest along interstate highways; McCaw had a similar aggregation in the Northeast. Both resulted in a spidery allocation, strange in appearance until it was compared with an interstate highway map. Moreover, new services might require spectrum in unanticipated geographic concentrations. It seemed important to let the market choose the allocation in the process of bidding.

Most of the academic economists involved in the design considered revenue an important goal. With the United States' substantial government debt and prevailing deficit, revenue had value beyond the strict dollar amount raised because increased revenue permitted a reduction in distorting taxation. Revenue has an additional importance, however. Revenue is evidence that high-value bidders are being selected; that is, revenue is itself evidence that the price system is working. In principle, of course, higher revenue need not signal allocative efficiency, as is easily seen by noting that the monopoly quantity has higher revenue than the efficient competitive price.

The 1993 legislation authorizing auctions directed the FCC to "disseminate licenses to a wide variety of applicants, including small businesses, rural telephone companies, and businesses owned by members of minority groups and women." Collectively these groups came to be known as "designated entities." Historically, the U.S. government has favored such groups by set-asides. In a set-aside a portion of the items for sale (or contracts to buy) is reserved for the designated entities. This is a potentially costly means of meeting the congressional requirement; McAfee and McMillan (1989) have argued that price preferences, in which the designated entity obtains a bidding credit or can submit a given bid at less cost than other bidders, are a more effective and less costly means of achieving the same ends. The FCC has used bidding credits, spectrum set-asides, and installment payments to promote participation by designated entities. It discontinued use of installment payments in 1997.

Most of the auction-design discussions among the economists, FCC staff, and participating companies came down to eight basic choices discussed in the subsequent sections.

Ascending versus Sealed Bids

The most fundamental choice in auction design is whether to use a once-and-for-all bid or to permit revision of bids over time. Governments generally choose sealed bids over ascending bids; corporations are more likely to prefer ascending bids.

The usual view is that a simultaneous sealed bid is more difficult to rig than an ascending auction. A member of a conspiracy to rig the bids can cheat on the conspiracy by submitting a bid secretly; thus sealed bids encourage breakdown of cartels. In contrast, a cartel can punish a deviation in an oral auction immediately. On the other hand, sealed bids open the door to collusion with a government official, who secretly reports on others' bids to one of the bidders. Sealed bids are less susceptible to collusion among the bidders but more susceptible to malfeasance by government officials. In any case, organized conspiracy was not anticipated to be a problem.

Ascending auctions differ from sealed bids in the release of information: in an ascending auction bidders can respond to the behavior of others during the course of the auction. This ability to revise bids in light of the behavior of others has consequences developed by Paul Milgrom and Robert Weber (1982) and on average increases the revenue in the auction in a symmetric environment. An ascending auction reveals information about the bidders to the bidders. Revealing information reduces the size of the information rents obtained by bidders and increases prices on average. This positive relationship between revenues and information transmission during the auction has been labeled "the linkage principle" in, for example, Milgrom (1989, p. 16).

In addition, when bidders are asymmetric, sealed-bid auctions tend to select inefficiently, while ascending auctions tend to be more efficient. Intuitively, the reason is that in a sealed-bid auction a strong bidder faces less competition than a weak bidder. The strong bidder faces the field minus itself, while the weak bidder faces a field that includes the strong bidder. Thus a strong bidder will seek a higher profit, while a weak bidder, facing stronger competition, will bid closer to actual value. The strong bidder will sometimes lose in circumstances when it has the higher value. In an ascending auction the strong bidder would revise its bid upward to win if it were actually the high-value bidder.

Ascending auctions have the virtue of reducing regret: bidders need not "leave money on the table," the euphemism for a bid substantially higher than the second-highest bid. Such bids make the bidder appear to pay too much, which could harm the bidder's career even if the bid was in fact ex ante optimal. In an ascending auction actual competition, rather than an expectation about competition, forces bids to the level achieved, and this feature reduces regret.

Although the weight of the economic literature favored ascending auctions, the case was far from transparent, primarily because of the history of government auctions using sealed bids, including the Department of the Interior's offshore oil auctions, which use a simultaneous sealed-bid auction of as many as 150 distinct tracts. That the Department of the Interior continues to use such a flawed mechanism demonstrates the difficulty of changing the status quo.

Simultaneous versus Sequential

Auctioneers Christie's and Sotheby's have sold billions of dollars of goods by sequential oral auctions, and there was a strong sense that what is good enough for antiques sellers is good enough for the U.S. government. Several economists persisted in the view that the simultaneous designs were unnecessarily complicated, both for the government to implement and for bidders to participate. This view still prevails in some circles. As will become clear, the nature of spectrum auctions renders this view naïve.

In a sequential design, licenses are ordered and then sold in a series of auctions. When the amount a bidder is willing to pay for an item depends on the other items it acquires, sequential auctions deny the bidder crucial information. Consider a bidder who values two items separately at $1 each, but collectively at $3. This bidder on the first item would be willing to pay up to $2, provided it expected the third item to sell for no more than $1. On the other hand, having bought the first item for $2, the bidder would now be willing to pay up to $2 for the second item, even though this creates a $1 loss on the pair. This problem for the bidder is known as the exposure problem: holding one item exposes one to a loss created by the complementarities in values. The bidder has to forecast the price of future items to bid sensibly on the earlier items. This need to forecast creates a dilemma for the bidder, whether to bid safely and probably lose, or bid aggressively and wind up stuck with an incomplete aggregation. Only a combinatorial design avoids the exposure problem, but at the expense of creating other problems.

Sequential auctions are problematic when items are substitutes, as well as complements. Suppose that two identical items are auctioned sequentially and a bidder wants only one. In bidding on the first item, the bidder must guess the selling price of the second item. There is no reason to expect that the selling price will be the same for both items, and bad guesses may result in the items not being awarded to the parties who value them the most. Empirical auction data have shown that this inefficiency is a problem. In particular, Gandal (1997) demonstrated that use of sequential auctions for cable television franchises in Israel affected revenues and arguably affected efficiency. Similarly, it has been argued that sequential auctioning by country of licenses intended for third generation (3G) cellular technology in Europe led to the inflation of early prices and collapse of prices in countries later in the sequence, although many economists view this as implausible.

A major problem for sequential designs is the order in which licenses are to be sold, especially when different orders advantage different firms. The orders of sequential sales of broadband PCS licenses proposed by different parties included largest to smallest, smallest to largest, east to west, and west to east. Ordering induces a bias; particular firms care more about some markets than others and will prefer learning from the markets they care less about first. In addition, sorting out the most important markets before less important complementary markets is advantageous. No ordering is neutral, which created a problem for supporters of sequential designs: they disagreed about the ordering. Thus there was little support for any one sequential design even though a majority of the economists initially supported sequential designs.

Another serious problem of sequential designs is timing. Proponents of sequential designs generally considered selling dozens of licenses per day. Because some of the most valuable licenses were expected to—and did—sell for hundreds of millions of dollars, a sequential design required making billion-dollar decisions in the course of a single day. Although companies can execute a billion-dollar decision in a day, they cannot reasonably incorporate new information in the process because changing a billion-dollar decision usually requires days or even weeks of deliberation at the senior management level. Running sequential auctions in a timely manner—completing hundreds in a month or two—would generally reduce the utility of information gleaned from the auctions themselves, forcing bidders into a prespecified set of strategies. This reduction vitiates the major advantage of ascending auctions: information release.

In contrast to sequential auctions, simultaneous designs were not well understood, and proponents had to solve a variety of problems. The most famous simultaneous design, the "silent auction" used by charities, has physical locations for each of the items for sale. Because bidding closes on all items at the same time, this design requires bidders to have as many bidding agents as the number of items they desire. The alternative, simultaneous multiple-round bidding (SMR), was proposed as a solution to this problem. In SMR all the licenses are available for bidding in each of discrete rounds. The minimum bid on any one license is increased over the previous round by a percentage. Bidders can bid on any or all the licenses in a given round, and the maximum bids from the previous round become the basis for the next round of bidding. This solution created a different problem. A bidder would like to see first what others are willing to pay and then choose the items that represent the best value. That is, every bidder has an optimal strategy of waiting until the others have completed their bidding. Were all the bidders to follow this strategy, of course, there would be no bidding at all. The standard solution in the literature, due to Charles Plott (see, e.g., Plott, Wit, and Yang, 2003), involves a positive probability of ending the auction, which itself creates a probability of inefficiency, regret, and frustration by participants and government alike.

The solution invented by Paul Milgrom and Robert Wilson is the activity rule—each bidder must be active (have the standing high bid or make a new bid) in every round on a specified fraction of the licenses it hopes to win. Thus a bidder that seeks twelve licenses would, under a 50% activity rule, be required to bid on at least six licenses. (After the first auction activity was defined in terms of bidding units to account for differences in license sizes. Each license was assigned bidding units equal to the product of the amount of spectrum in MHz and population in the license area.) This puts pressure on bidders to bid but still allows them the flexibility to substitute to other licenses if the licenses on which they are currently bidding become too expensive.

The magnitude of the activity rule determines the ease of substitution. With a low activity requirement, bidders can readily substitute from one set of licenses to another because they hold few of the licenses they hope to purchase and thus have lots of "purchasing capacity." If bidders are generally bidding only on the licenses required to maintain activity, price increases will stall when the potential demand for licenses at current prices is about one over the amount of the activity rule times the available amount of spectrum. Thus the activity rule needs to be tightened over the course of the auction, or prices will come to a stop even though excess demand remains. The first auctions used a three-phase system: initially a 33% activity requirement, then a 67% requirement, followed by a 100% requirement.[5]

This design, created primarily by academic economists, has become known as the FCC auction. The key advantages of this auction design relative to all others are that it promotes (1) the timely release of information and (2) substitution by buyers. If a price increase on one or more licenses renders that group more expensive, a buyer can switch to another group, usually with no penalty but at worst with a modest penalty. This substitution is facilitated by several complementary features, including the activity rule itself, which directly permits substitution, the gradual tightening of the activity rule,[6] the design of the licenses for sale into "similar-sized" units, and the modest withdrawal penalty that forces bids to be meaningful expressions of willingness to pay while still encouraging the assembly of efficient aggregations.

The success of the FCC auction is perhaps most dramatically illustrated by the second narrowband (paging-frequency) licenses. Here thirty licenses were sold. Six frequency bands were sold in each of five geographic regions. Four of the sets of regional licenses were aggregated into national licenses, and similar licenses sold for similar prices; indeed, the single price anomaly was the consequence of a failed attempt to game the auction, resulting in a withdrawal and a low price for that license. (With the withdrawal penalty, the license sold for a similar price.) Bidders were able to aggregate the licenses they desired, and prices appeared competitive.

Combinatorial Bids

A combinatorial bid involves the ability to bid on a package of licenses as an all-or-nothing bid. Various combinatorial bidding methods were promoted by academic economists. The two main forms initially proposed were (1) to permit bidding on any subset and allow the bidders to revise their bids, using an SMR format, until bidding stopped, or (2) have individual licenses plus a national license.

Combinatorial auctions have the advantage that bidders can express their demands directly—a bidder who wants three licenses can bid on them as a unit. Combinatorial auctions provide a solution to the exposure problem identified earlier. However, combinatorial auctions are more complex because there are many potential combinations. There is also the "threshold" or "free-rider" problem—the potential that bidders on subsets of a larger package will not be able to combine their bids to beat the bid on the larger package even when the sum of the value of the subsets exceeds the value of the package. As with all free-rider problems, there is a risk that the bidders will fail to bid enough. The risk of lower revenue and inefficiency due to the threshold problem, along with the potential complexity of implementing a combinatorial auction, deterred the FCC from selecting it in spite of serious support among academic economists. Nevertheless, combinatorial auctions remain a vibrant area of research and a continuing interest of the FCC. Indeed, since 1996, significant experimental and theoretical research has suggested ways to mitigate the threshold problem (see, e.g., DeMartini et al. 1998; Marx and Matthews 2000; and Ausubel, Cramton, and Milgrom 2006).

Reserve Prices

In the traditional theory of auctions, reserve prices represent a trade-off between revenue and efficiency; efficient reserve prices are the seller's value, usually set to zero as a modeling artifact, while the monopoly reserve price is substantially in excess of the seller's value. Economists divided on the issue of reserve prices along the lines of the importance they gave to efficiency and revenue.

From a static point of view, any positive reserve price conflicts with efficiency when a license is not sold.[7] However, a new, dynamic view of reserve prices emerged from the FCC auctions. A seller who fails to sell has the option of selling in the future; thus the seller's value may not be the seller's use value but instead the value of a buyer not yet present. If the seller expects to sell at a higher price in the future, it may be efficient not to sell today. This may be the case when new technologies are possible, technologies that will be resisted by incumbents because they supplant existing services. The question then is whether private parties or the FCC should keep this "option value." In light of allegations of stockpiling of spectrum by incumbents, reserving spectrum for future technologies could be important. The threat of stockpiling by incumbents provides a rationale for public rather private holding of idle spectrum,[8] but there is by no means a consensus on this point.

Size of Increments

The use of rounds of bidding forces a decision on a bid increment—what is the minimum increase between rounds? If it is too small, the auction will take too many rounds, wasting time and effort. If it is too large, the auction will allocate the licenses inefficiently. The consensus recommendation was created by an insight of Paul Milgrom. Halving the step size or increment approximately doubles the time required to complete the auction. However, inefficiency arises only when the bidder with the second-highest value wins because the high-value bidder is unwilling to pay the second-highest value plus an increment. The loss in value is at most an increment and, moreover, occurs only when the second-highest value is close to the highest value, which has a probability proportional to the size of the increment, so the loss is approximately proportional to the bid increment squared. Consequently, even quite large increments—say, 10%—have an efficiency impact proportional to the square of the increment, or 1%. Milgrom's argument persuaded most economists that substantial increments were not very costly to efficiency and hence represented the best way to speed up auctions. The exception to the consensus was, by and large, from those in favor of continuous-time auctions.

In spite of the consensus among economists, business and the FCC remained leery of large increments. The FCC instead employed an increment accelerator rule, which determined the size of the increment by the amount of competition present in preceding rounds.

Bidder Preferences

An innovative feature of the FCC auction was the use of price preferences or bidder credits, in addition to set-asides, to favor designated entities. A bidder credit is akin to the idea of handicapping in sports. Suppose that a designated entity, because of lack of capital or expertise, has a 10% lower value on average. Then the auction can level the competition by providing a 10% bidder credit—that is, charging the designated bidder only 90% of its bid.

Bidder credits are advantageous over set-asides for at least seven distinct reasons. First, bidder credits rather than set-asides increase competition in an SMR auction. The designated entities become more effective bidders thanks to the handicapping, while the nondesignated entities are some competition for the licenses that would otherwise be set aside. Second, bidding credits set a price or value for promotion of designated entities, thereby permitting resale at the cost of refunding the bidder credit.[9] Third, the inefficiency of promoting the designated entities may be reduced because designated entities win auctions only when they are close to the willingness to pay of the other firms, relative to a set-aside. Fourth, revenues may increase over the levels that would prevail without any bidder credits because for small levels of bidder credit, the increase in competition outweighs the inefficient allocation.[10] Of course, revenue will fall as the bidder credit becomes larger than the disadvantage of the designated group. Fifth, the tendency for set-asides is to create very small licenses, so-called ghetto licenses, perhaps because nondesignated entities lobby hard to prevent the setting aside of valuable licenses that could lead to meaningful competition. Thus a universal modest bidder credit, applicable to all licenses, will prevent the creation of ghetto licenses and permit the designated entities to compete effectively for all licenses. Sixth, price preferences are a versatile instrument, capable of giving distinct levels of preference to different categories of designated entities; for example, the credit may vary with the size of the small business, although the Supreme Court has ruled price preferences like bidder credits for telecommunications auctions illegal.[11] Seventh, price preferences naturally apply to partial ownership by giving partial credit.

The FCC used 25% bidder credits in the first narrowband auction, but none of the licenses were won by designated entities. The credit level was increased to 40% in the second narrowband auction, and all the licenses available for bidder credit were won by designated entities. It is not clear, unfortunately, whether the success in the second auction was due to the increased bidder credit or the increased time to prepare; some of the designated entities seem to be firms that came into existence to take advantage of the bidder credits. Indeed, supporting the "designer-firm" theory, the actual prices paid for the licenses net of the bidder credit were just slightly below the prices paid for licenses not eligible for bidder credits, suggesting that the designated entities were nearly competitive with the nondesignated entities.

Spectrum Caps

Spectrum caps—a limit on the amount of spectrum any one participant may hold in a given geographic area—are a simple means of ensuring adequate competition in the final product market. The existing structure of the wireless market was a duopoly by design. The sale was intended to foster competition, and the spectrum caps were set such that an incumbent could not purchase a 30 MHz license in its existing footprint; this guaranteed at least three new entrants in every market. The FCC aimed to have five competing firms in every geographic market.[12] An important characteristic of spectrum caps is that they are substantially easier to enforce than antitrust laws.

Revealing Bidder Identity

In a standard oral auction it is difficult, although not impossible, to conceal the identities of the bidders. However, in an auction with multiple rounds of bidding done on computers, it becomes possible to conceal the identity of the bidders from other bidders. Concealment can be accomplished by reporting only the price bid on each license and not the identity of the firm that bid, or by giving bidders identification numbers and revealing the bidder number but not the mapping from real bidders to bidder identifications.

The disadvantage of concealing bidder identities is that there is information content not just in what was bid but in who did the bidding. In particular, a bid from a company completing a national license by bidding on a "hole" may be a less significant indicator of value than a bid by a start-up. Perhaps more empirically relevant, there are technological economies of scale, so knowing whether a large Global System for Mobile communications (GSM) technology carrier is bidding may matter substantially to the smaller carriers, who hope to have access to handsets that will become available only if a large player buys spectrum. Thus the identity of other bidders could potentially affect valuation. A second disadvantage, which was entertaining if not economically significant in the first narrowband auction and potentially significant in later auctions, is that bidders spend a lot of time trying to figure out which bidder identification corresponds to which bidder.

There are two main advantages to reporting only prices. First, because concealed identities make defection from a cartel agreement harder to punish, concealing identities should reduce the likelihood of behavior ranging from price-fixing to "tacit understandings." Second, concealed identities make it much harder to "game" the auction. An example of such gaming occurred in the Australian spectrum auctions in which one bidder was attempting to assemble five adjoining licenses and the spoiler continued to bid on the middle license in a strategy that became known as "giving the middle finger." When all that is observable is the price, such a punishment strategy is much less likely to succeed. For a description of the use of such strategies in an FCC auction, see Cramton and Schwartz (2002). In the end, the FCC chose to conceal identities but to reveal bidder identifications in the first auction. Because this auction was run by bidding in a booth like a polling booth, it was relatively easy to deduce which bidder identification went with which company.

How to Design Efficient and High-Revenue Auctions

How well did the economists do in designing the auctions, and what have we learned since? The FCC has run over seventy auctions since 1994. From this wealth of experience we can distill several key lessons and principles. In summary form, these principles are the following:

- Promote information revelation:
  Use ascending auctions.
  Provide bidding tools for analyzing bids.
- Promote substitution:
  Sell licenses simultaneously.
  Make the licenses similarly sized.
  Squash and "round" license sizes for the auction activity-rule calculations.
  Progressively tighten activity rules.
- Make the problem to be solved by participants easier:
  Use fixed increments.
  Conceal identities of bidders.
  Allow enough time between rounds for good decisions.
- Design simple, transparent rules.

Conclusion

Academic economists, working with FCC staff and especially FCC economist Evan Kwerel, created an auction form that has been used to sell over

$100 billion dollars of spectrum in dozens of countries. The design reflected trade-offs that were understood only because of the development of auction theory in the 1980s and thus implemented recent innovations in economic analysis. The FCC auction performed well by a variety of measures and seems to have balanced revenue and efficiency.

The FCC auction has rightfully been credited as an important impetus to the adoption of game-theoretic techniques in the business world. The usefulness of economists in the private sector, long a subject of derision, was forcefully demonstrated by the application of sophisticated game theory in this real-world setting.

The role of economics and economists in the FCC auctions was not to prove a theorem that provided an optimal mechanism for some particular setting. Instead, it was to use existing theory to identify the most important factors, design a system that accommodates the needs of market participants, generally in as simple a manner as possible, and identify the features that may be important. Theory, then, is a guide, not a destination.

Notes

John McMillan passed away before this chapter was complete, so it does not reflect his gift for engaging exposition. This chapter does not necessarily represent the views of our employers.

1. *Omnibus Reconciliation Act of 1993,* U.S. Public Law 103-66, 103rd Cong., 1st sess., August 10, 1993.
2. Section 309(j) of the *Communications Act,* 47 U.S.C. § 309(j).
3. A Vickrey (or second-price) auction is a sealed-bid auction in which the high bidder winning but paying the second highest bid. See McAfee and McMillan (1987 p. 710). It is equilibrium behavior only in the relatively specialized private-values environment.
4. MCI, the only major telephone company with no existing spectrum holdings, was a strong proponent of the national license. This was one of the most obvious examples of self-interested lobbying.
5. A literal 100% requirement, when applied to unequally sized licenses, is bad because of integer problems. Integer problems themselves gave bidders headaches in FCC auctions. The Mexican spectrum auctions had the additional problem that Mexico City was a huge portion of the entire country. As a result, the modern design "squashes" or levels the license size, preferably to values that are easy to combine into substitutes. In Mexico the license "sizes," for the purposes of evaluating the activity rule, were set at 1, 2, 3, 4, 6, 12 (Guadalajara), and 24 (Mexico City) to promote substitution. These are quite different from population sizes; the squashing makes it easier to substitute between Mexico City and a group of other licenses.

6. A 100% activity rule makes substitution difficult. Suppose that a bidder seeks six contiguous licenses but does not care where the six are. At 100% a bidder can move without risk of penalty only on the licenses whose prices increased; if prices of three increased, the bidder faces a substantial risk of being left either paying penalties or stuck with the remaining three.
7. Unused spectrum is lost. Use of the spectrum does not inhibit future use of the spectrum. Thus social efficiency calls for full use. This full use is mitigated by an installed base of equipment to use the spectrum, and the sale of spectrum interacts with standards on use.
8. See Cramton, Skzyracz, and Wilson (2007).
9. The amount of the bidder credit should not be confused with the social value of favoring the designated entities or with the size of the disadvantage they face. Generally, optimal bidder credits will be somewhat less than the social value of favoring the entities because the bidder credit must trade off the revenue obtained against the desire to increase participation by the entities. Second, the bidder credit may be larger or smaller than the disadvantage faced by the entities. Revenue maximization requires bidder credits that are positive but smaller than the disadvantage because fully compensating for the disadvantage causes the designated entities to win too often. See McAfee and McMillan (1996, p. 167).
10. Allocative inefficiency has a zero first-order effect, while the change in competition has a first-order effect, when the credit is small.
11. *Adarand v. Pena,* 15 S. Ct. 2097 (1995).
12. The FCC subsequently abolished the ownership cap and drafted industry-specific merger review guidelines that were never released.

References

Ausubel, Larry, Peter Cramton, R. Preston McAfee, and John McMillan. "Synergies in Wireless Telephony." *Journal of Economics and Management Strategy,* 6, no. 3 (Fall 1997), 497–527.

Ausubel, Larry, P. Cramton, and P. Milgrom. "The Clock-Proxy Auction: A Practical Combinatorial Auction Design." Chapter 5 in Peter Cramton, Yoav Shoham, and Richard Steinberg, eds., *Combinatorial Auctions.* Cambridge, Mass.: MIT Press, 2006.

Coase, Ronald H. "The Federal Communications Commission." *Journal of Law and Economics,* 2 (October 1959), 1–40.

Cramton, Peter, and Jess Schwartz. "Collusive Bidding in the FCC Spectrum Auctions." *Contributions to Economic Analysis & Policy,* 1, no. 1 (2002).

Cramton, Peter, Andrzej Skrzypacz, and Robert B. Wilson. Available from www.cramton.umd.edu/papers2005-2009/cramton-skrzypacz-wilson-competition-in-700-mhz-auction.pdf.

DeMartini, Christine, Anthony M. Kwasnica, John O. Ledyard, and David Porter. "A New and Improved Design For Multi-object Iterative Auctions." HSS Working Paper 1054, California Institute of Technology, 1998.

Gandal, Neil. "Sequential Auctions of Interdependent Objects: Israeli Cable Television Licenses." *Journal of Industrial Economics,* 45 (September 1997), 227–244.

Marx, Leslie, and Steve Matthews. "Dynamic Voluntary Contribution to a Public Project." *Review of Economic Studies,* 67 (2000), 327–358.

McAfee, R. Preston, and John McMillan. "Auctions and Bidding." *Journal of Economic Literature,* 25, no. 2 (June 1987), 699–738.

———. "Government Procurement and International Trade." *Journal of International Economics,* 26 (1989), 291–308.

———. "Analyzing the Airwaves Auction." *Journal of Economic Perspectives,* 10 (Winter 1996), 159–175.

McMillan, John. "Selling Spectrum Rights." *Journal of Economic Perspectives,* 8, no. 3 (Summer 1994), 145–162.

Milgrom, Paul. "Auctions and Bidding: A Primer." *Journal of Economic Perspectives,* 3 (Summer 1989), 3–22.

———. *Putting Auction Theory to Work*. Cambridge: Cambridge University Press, 2004.

Milgrom, Paul, and Robert Weber. "A Theory of Auctions and Competitive Bidding." *Econometrica,* 50, no. 5 (September 1982), 1089–1122.

Plott, Charles, J. Wit, and W. C. Yang. "Parimutuel Betting Markets as Information Aggregation Devices: Experimental Results." *Economic Theory,* 22 (2003), 311–351.

Safire, William. "Essay; The Greatest Auction Ever." *New York Times*, Thursday, March 16, 1995, Section *A*, p. 25.

COMMENT

Jeremy Bulow

Why was the simultaneous multiple-round (SMR) auction design so successful in the 1995 FCC AB auction and in many other settings? Ironically, one reason is that its structure almost guaranteed the sequential closure of markets, with the biggest and most valuable markets closing first. Activity rules made it easier to move from larger licenses to smaller ones in the late bidding rather than the other way around. This probably raised revenue and helped with problems of complementarity.

To see the revenue point, consider Example 1, in which there are two licenses and three bidders with the following valuations.

Example 1

	License A	License B
Bidder 1	10	20
Bidder 2	20	40
Bidder 3	30	60

Say each bidder is allowed to buy only one license. If the licenses are sold sequentially, with the smaller A sold first, prices will be 10 for license A (won by bidder 2) and 20 for license B (won by 3, competing only against 1). If B is sold first, then the prices will be 30 for license B (sold to 3 competing against 2, whose alternative is to win license A at a price of 10), and

10 for A. Auction rules that make it much easier to move from bigger licenses to smaller ones than to go the other way often effectively assure that smaller licenses will not clear before the bigger ones.[1]

Having larger licenses clear first provides a further advantage for aggregation. In packing a car trunk, most people put the big suitcases in first. Similarly, in forming an aggregation, most bidders will focus on the largest, most valuable licenses in a package. A bidder interested in greater New York may regard a partial aggregation that includes the city but not all suburbs as a success and one that contains all the suburbs but not the city as a disastrous failure.[2]

A second reason to like the SMR is that it has some significant advantages over a pure Vickrey auction, in which each bidder writes down its value for all license combinations and a computer calculates the value-maximizing allocation, each bidder receiving a surplus equal to its marginal contribution to total value. The authors discuss how multiple-round bidding allows bidders to learn information through the auction, allowing them to bid more efficiently than in a one-shot auction. But there is a second disadvantage to a Vickrey approach.[3]

Assume that bidders have the valuations listed in Example 2.

Example 2

Bidder	Value of A only	Value of B only	Value of AB
1	0	0	10
2	9	0	9
3	0	9	9

The Vickrey auction will award the licenses to bidders 2 and 3 at prices of 1 each. (Since the total value in the auction is 18, and it would be only 10 without either 2 or 3 they each receive a surplus of 8). This means that if these bidders each have actual values of 6 instead of 9 for their licenses, they can gain by colluding by raising their bids to 9 each (lowering their prices from 4 to 1), and if they have values of 2 each, they can still profit by bidding 9 on the licenses. The Vickrey auction is thus deeply susceptible to collusion, because jointly raised bids can allow the colluders to acquire more licenses at lower prices. Furthermore, the "collusion" can take the form of one entity submitting bids as two—if bidders 2 and 3 were treated as one bidder, they would have to pay 10 to buy the licenses.

Of course all mechanisms will fail if bidders can collude. The difference is that in most economic models, for collusion to be successful it must be "collusion in the large." The competitive model works well because in most settings collusion requires the participation of almost all bidders. With Vickrey-Clark-Groves mechanisms one may only require "collusion in the small"; given the top bid for a package, two players bidding on individual licenses may be able to profitably collude on their own. (In fact, the bids in Example 2 constitute a Nash Equilibrium even with low bidder values, although collusion would be required under some refinements.)[4]

Finally, even when the SMR is theoretically optimal, as in European 3G spectrum auctions, where each bidder was limited to one license, successes (the United Kingdom) and failures (Switzerland, Italy) depended on whether the number of licenses and other ancillary rules encouraged entry and deterred collusion. The ultimate goal of any auction design, which the SMR often but not always achieves, must be to facilitate vigorous and relatively efficient bidder competition. The SMR's designers performed a valuable service for consumers and taxpayers around the world, but as with any prescription, it must be used with care.

Notes

1. Of course, the key is that the gaps in value are larger for the big licenses, not the absolute values, but this should usually be true.
2. Some of the largest aggregation problems are solved by the initial specification of licenses: if all agree that New York City should be sold to one bidder then the auction will likely bundle the boroughs into one license.
3. There are other disadvantages to the Vickrey auction, for example when bidders have common values or budget constraints or are risk averse.
4. I originally raised these points about the Vickrey auction in a presentation at the FCC Combinatorial Bidding Conference in Wye River, MD, May 5,2000. The presentation was based on discussions with Paul Klemperer.

—-1
—0
—+1