Putting Auction Theory to Work: The Simultaneous Ascending Auction

Paul Milgrom

Stanford University

I review the uses of economic theory in the initial design and later improvement of the "simultaneous ascending auction," which was developed initially for the sale of radio spectrum licenses in the United States. I analyze some capabilities and limitations of the auction, the roles of various detailed rules, the possibilities for introducing combinatorial bidding, and some considerations in adapting the auction for sales in which revenue, rather than efficiency, is the primary goal.

I. Introduction

The "simultaneous ascending auction" was first introduced in 1994 to sell licenses to use bands of radio spectrum in the United States. Much of the attention devoted to the auction came from its role in reducing federal regulation of the radio spectrum and allowing market values, rather than administrative fiat, to determine who would use the spectrum resource. Many observers were also fascinated by the then-novel use of weblike interfaces for bidders. The large amounts of money involved were yet another source of interest. The very first use of the auction rules was a $617 million sale of 10 paging licenses in July 1994. In the broadband personal communications services (PCS) auction, which began in December 1994, 99 licenses were sold for a total price of approximately $7 billion. Once the auctions had been conducted, it became much harder to ignore the tremendous value of the large amounts of spectrum allocated to uses

My thanks go to Peter Cramton, Paul Klemperer, Yoav Shoham, and Padmanabhan Srinagesh, as well as seminar participants at Stanford, the University of Pittsburgh, and Yale for comments on an earlier draft.

[*Journal of Political Economy*, 2000, vol. 108, no. 2]

such as high-definition television, for which Congress had demanded no compensation at all. Moreover, the perceived successes with the new rules inspired imitators to conduct similar spectrum auctions in various countries around the world and to recommend similar auctions for other applications.

Among academic economists, attention was also piqued because the auction design made detailed use of the ideas of economic theory and the recommendations of economic theorists. Indeed, the U.S. communications regulator adopted nearly all its important rules[1] from two detailed proposals for a simultaneous ascending auction: one by Preston McAfee and the other by Robert Wilson and me. Economic analysis dictated nearly all the rule choices in the first few auctions. Various reviews suggest that the new auction design realized at least some of the theoretical advantages that had been claimed for it.[2]

Several parts of economic theory proved helpful in designing the rules for the simultaneous ascending auction and in thinking about how the design might be improved and adapted for new applications. After briefly reviewing the major rules of the auction in Section II, I turn in Section III to an analysis based on *tatonnement* theory, which regards the auction as a mechanism for discovering an efficient allocation and its supporting prices. The analysis reveals a fundamental difference between situations in which the licenses are mutual substitutes and others in which the same licenses are sometimes substitutes and sometimes complements. When the licenses are mutual substitutes for all bidders, not only is it true that equilibrium prices exist, but straightforward, ''myopic'' bidding in the auction leads bidders to prices and an allocation that are close to competitive equilibrium. This happens even though, in contrast to traditional *tatonnement* processes, prices in the auction process can never fall and can rise only by fixed increments. However, if even one bidder has demand in which licenses are not mutual substitutes, then there is a profile of demands for the other bidders, all of which specify that licenses are mutual substitutes, such that no competitive equilibrium prices exist. There is an inherent limitation in the very conception of the auction as a process for discovering a competitive allocation and competitive prices in that case.

Section IV is a selective account of some applications of game theory to evaluating the design of the simultaneous ascending auction for spectrum sales. Game-theoretic arguments were among those

[1] The sole exceptions were the financing rules, which were devised to encourage participation in the auctions by financially weak smaller businesses and those owned by women and minorities.

[2] See Cramton (1995), Milgrom (1995), and McAfee and McMillan (1996) for accounts of the auction and the run-up to it.

that convinced regulators to adopt my suggestion of an "activity rule," which helps ensure that auctions end in a reasonable amount of time. Game theory also provided the decisive argument against the first "combinatorial bidding" proposals and has also been employed to evaluate various other suggested rule changes.

Results like those reported in Section III have led to renewed interest in auctions in which bids for license packages are permitted. In Section V, I use game theory to analyze the biases in a leading proposal for dynamic combinatorial bidding. Section VI briefly answers two additional questions that economists often ask about auction design: If trading of licenses after the auction is allowed, why does the auction form matter at all for promoting efficient license assignments? If the number of licenses to be sold is held fixed, how sharp is the conflict between the objectives of assigning licenses efficiently and obtaining maximum revenue? Section VII presents a conclusion.

II. Simultaneous Ascending Auction Rules in Brief

A simultaneous ascending auction is an auction for multiple items in which bidding occurs in rounds. At each round, bidders simultaneously make sealed bids for any items in which they are interested. After the bidding, round results are posted. For each item, these results consist of the identities of the new bids and bidders[3] as well as the "standing high bid" and the corresponding bidder. The initial standing high bid for each item is given (it may be zero), and the "corresponding bidder" is the auctioneer. As the auction progresses, the new standing high bid at the end of a round for an item is the larger of the previous standing high bid or the highest new bid, and the corresponding bidder is the one who made that bid. In addition to the round results, the minimum bids for the next round are also posted. These bids are computed from the "standing high bid" by adding a predetermined bid increment. For spectrum licenses, the increments are typically the larger of some fixed amount or a fixed percentage of the standing high bid.[4]

A bid represents a real commitment of resources by the bidder. In the most common version of the rules, a bidder is permitted to

[3] The first trial of the simultaneous ascending auction did not include announcements of bidder identities, but the larger bidders were often able to infer identities anyway, leading to a change in the rules to remove that advantage. The practice of identifying the bidders continues to be controversial.

[4] In the spectrum auctions, the percentage has usually been 5 percent or 10 percent (and in recent auctions has been dependent on the level of bidding in the auction). The appropriate size of the increment has also been subjected to economic analysis that takes into account the cost of adding rounds to the auction and the extent and type of the uncertainty about bidder values.

withdraw bids, but there is a penalty for doing so: if the selling price of the item is less than the withdrawn bid, the withdrawing bidder must pay the difference. In other applications, bid withdrawals are simply not permitted.

A bidder's eligibility to make new bids during the auction is controlled by the "activity rule." Formally, the rule is based on a "quantity" index, such as spectrum bandwidth or population covered by a license, that roughly corresponds to the value of the license. During the auction, a bidder may not have active bids on licenses that exceed its *eligibility*, measured in terms of the index.

At the outset of the auction, each bidder establishes its *initial eligibility* for bidding by making deposits covering the quantity of spectrum for which it wishes to be eligible. During the auction, a bidder is considered *active* for a license at a round if it makes an eligible new bid for the license or if it owns the standing high bid from the previous round. At each round, a bidder's activity is constrained not to exceed its eligibility. If a bid is submitted that exceeds the bidder's eligibility, the bid is simply rejected.

The auction is conducted in a sequence of three stages, each consisting of multiple rounds. The auction begins in stage 1, and the administrator advances the auction to stage 2 and later to stage 3 when there are two or more consecutive rounds with little new bidding. In each round during stage j, a bidder that wishes to maintain its eligibility must be active on licenses covering a fraction f_j of its eligibility. If a bidder with eligibility x is active on a license quantity $y < f_j x$ during stage j, then its eligibility is reduced at the next round to y/f_j.[5]

The activity rules have two functions. First, they create pressure on bidders to bid actively, which increases the pace of the auction. Second, they increase the information available to bidders during the auction, particularly late in the auction. For example, in stage 3, bidders know that the remaining eligible demand for licenses at the current prices is just $1/f_3$ times the current activity level, which can be rather informative when f_3 is close to one.[6]

The auction also provides five "waivers" of the activity rule for each bidder, which can be used at any time during the auction, that allow the bidder to avoid reduction in its eligibility in a given round. The waivers were included to prevent errors in the bid submission process from causing unintended reductions in a bidder's eligibility, but they also have some strategic uses.

[5] In the 1998 auction of licenses to spectrum in the 220 MHz range, the fractions used were $(f_1, f_2, f_3) = (.8, .90, .98)$.

[6] See Cramton (1997) for evidence on the informational content of bids.

There are several different options for rules to close the bidding that were filed with the regulator. One proposal, made by Preston McAfee, specified that when a license had received no new bids for a fixed number of rounds, bidding on that license would close. That proposal was coupled with a suggestion that the bid increments for licenses should reflect the bidding activity on a license. A second proposal, made by Robert Wilson and me, specified that bidding on all licenses should close simultaneously when there is no new bidding on any license. To date, the latter rule is the only one that has been used.[7]

When the auction closes, the licenses are sold at prices equal to the standing high bids to the corresponding bidders. The rules that govern deposits, payment terms, and so on are quite important to the success of the auction,[8] but they are mostly separable from the other auction rule issues and receive no further comment here.

III. Auctions and *Tatonnement* Theory

The simultaneous ascending auction is a process that, on its surface, bears a strong resemblance to the *tatonnement* process of classical economics. Like the *tatonnement* process, the objective of the auction is to identify allocations (which the spectrum regulators call "assignments") and supporting prices to approximate economic efficiency. Yet there are differences as well. First, bids in the auction represent real commitments of resources, and not tentatively proposed trades. Consequently, bidders are reluctant to commit themselves to purchases that may become unattractive when the prices of related licenses change. Second, in the auction, prices can never decrease. That is an important limitation because the ability of prices to adjust both upward and downward is a fundamental requirement in theoretical analyses of the *tatonnement*. Third, in the initial version of the simultaneous ascending auction, the bidders themselves name the prices. That contrasts with the Walrasian *tatonnement*, in which some

[7] The activity rule and the closing rule make this auction perform very differently from the otherwise similar "silent auction" commonly used in charity sales. In a silent auction, the items being sold are typically set on tables in a room and bidders walk around the room, entering their bids and bidder identification on a paper sheet in front of the items. Bidding closes at a predetermined time. It is a common experience that bidders in silent auctions often delay placing their bids until the final moment, completing their entry on the paper just as the bidding closes. With its closing and activity rules, the simultaneous ascending auction eliminates both the "final moment" that bidders exploit in silent auctions and also bidders' ability to wait until late in the auction before making any serious bids.

[8] Failure to establish these rules properly led to billions of dollars of bidder defaults in the U.S. "C-block auction." Similar problems on a smaller scale occurred in some Australian spectrum auctions.

fictitious auctioneer names the prices. Other differences arise from the nature of the application. The licenses sold in the auction are indivisible. This fact means that the set of allocations cannot be convex, so the usual theorems about the existence of competitive equilibrium do not apply. My analysis focuses on all these issues: the risk that bidders take when they commit resources in early rounds of the auction, the existence of competitive equilibrium, and whether the simultaneous auction process in which prices increase monotonically can converge to the equilibrium.

Let $L = \{1, \ldots, L\}$ be the set of *indivisible* licenses to be offered for sale. Denote a typical subset of L by S. In describing license demand, I also use S to represent the vector $\mathbf{1}_S$.

Assume that a typical bidder i who acquires the set of licenses S and pays an amount of money m for the privilege enjoys utility of $v_i(S) - m$. Given a vector of prices $\mathbf{p} \in \mathbb{R}^L_+$, $\mathbf{p} \cdot S$ denotes the total price of the licenses composing S. The demand correspondence for i is defined by $D_i(\mathbf{p}) = \text{argmax}_S\{v_i(S) - \mathbf{p} \cdot S\}$. Assume that there is *free disposal*, so $S \subset S'$ implies that $v_i(S) \leq v_i(S')$. I sometimes omit the subscript from demand functions, relying on the context to make the meaning clear.

During an auction, it often happens that a bidder is the high bidder on a subset of the licenses it would wish to acquire at the current bid prices. To describe such situations, I introduce the following definition: An individual bidder demands the set of licenses T at price vector $\mathbf{p}$, written $T \in X(\mathbf{p})$, if there exists $S \in D(\mathbf{p})$ such that $S \supset T$.

The usual definition of substitutes needs to be generalized slightly to deal with the case of demand correspondences. The idea is still the same, though: raising the prices of licenses not in any set S cannot reduce the demand for licenses in the set S.

DEFINITION. Licenses are *mutual substitutes* if for every pair of price vectors $\mathbf{p}' \geq \mathbf{p}$, $S \in X(\mathbf{p})$ implies that $S \in (X(p_S, p'_{L\setminus S}))$.

After any round of bidding, the minimum bids for the next round are given by the rule described in Section II. If the standing high bids at a round are given by the vector $\mathbf{p} \in \mathbb{R}^L_+$, then the minimum bid at the next round for the lth license is $p_l + \epsilon \max(p_l, \hat{p}_l)$ for some $\epsilon > 0$. The vector of minimum bids is then $\mathbf{p} + \epsilon(\mathbf{p} \vee \hat{\mathbf{p}})$, where $\hat{\mathbf{p}} \in \mathbb{R}^L_{++}$ is a parameter of the auction design, and the "join" $\mathbf{p} \vee \hat{\mathbf{p}}$ denotes the price vector that is the component-wise maximum of $\mathbf{p}$ and $\hat{\mathbf{p}}$.

During a simultaneous ascending auction, the minimum bid increment drives a wedge between the prices faced by different individual bidders. To analyze the progress of the auction, it is useful to define the *personalized* price vector $\mathbf{p}^j$ facing bidder j at the end of a round to be $\mathbf{p}^j = (p_{S_j}, (\mathbf{p} + \epsilon(\mathbf{p} \vee \hat{\mathbf{p}}))_{L\setminus S_j})$. That is, j's prices for the licenses S_j that it has been assigned are j's own standing high bids, but its prices for the other licenses are the standing high bids plus the minimum bid increment. This reflects the fact that under the rules of the auction, j can no longer purchase those other licenses at their current standing high bids.

My analysis of the *tatonnement* process consists of a study of what happens to bidder j when it (possibly) alone bids in a "straightforward" (nonstrategic) manner, and what happens when all bidders bid in a straightforward manner. When I say that j bids "straightforwardly," I mean that if, at the end of some round n, bidder j demands the licenses assigned to it (formally, if $S_j \in X_j(\mathbf{p}^j)$), then j makes the minimum bid at round $n + 1$ on a maximal set of licenses T such that $S_j \cup T \subset D_j(\mathbf{p}^j)$. Of course, bidders that wish to acquire multiple licenses commonly have an incentive to withhold some of their demand in order to reduce prices.[9] Consequently, the part of the analysis employing straightforward bidding must be understood as only a partial analysis, which ignores strategic incentives to highlight important nonstrategic properties of the auction design.

Intuitively, whenever the auction allows, the straightforward bidder bids to acquire the set of licenses that it demands at its personalized prices. Notice that the antecedent condition is automatically satisfied at the beginning of the auction because no bidder has yet been assigned any licenses.

Straightforward bidding often leads to ties at some rounds of the auctions. For the analysis of this section, any tie-breaking rule that selects a winner from among the high bidders will work.

My first theorem says that if j bids straightforwardly from the beginning of the auction and if licenses are mutual substitutes for j, then the antecedent condition for straightforward bidding continues to be satisfied round after round.

THEOREM 1. Assume that all the licenses are mutual substitutes for bidder j. Suppose that, at the end of round n, bidder j's assignment $S_j \in X_j(\mathbf{p}^j)$. If, at round $n + 1$, bidder j bids straightforwardly, then, regardless of the bids made by others, j's assignment S_j' at the end of round $n + 1$ satisfies $S_j' \in X_j(p^{j\prime})$, where $p^{j\prime}$ is j's personalized price at the end of round $n + 1$. Moreover, j's tentative profit after any round—what it would earn if the auction were terminated after

[9] Ausubel and Cramton (1996) argue that an incentive of this sort is unavoidable in a wide class of auctions, including all those that establish uniform prices for identical objects.

that round at the then-current prices and allocation—is always nonnegative.

Proof. Let T be the set of licenses on which j bids at round $n + 1$. Then, by the rules of the auction, $S'_j \subset S_j \cup T$. Also, by the hypothesis of straightforward bidding, $S_j \cup T \in D_j(\mathbf{p}^j)$. So, by definition, $S'_j \in X_j(\mathbf{p}^j)$.

The rules also imply that j's personalized prices $\mathbf{p}^{j\prime}$ for the licenses in S'_j coincide with the prices of those licenses according to $\mathbf{p}^j$. Moreover, j's personalized prices cannot fall from round to round: $\mathbf{p}^{j\prime} \geq \mathbf{p}^j$. Hence, by the definition of mutual substitutes, $S'_j \in X_j(\mathbf{p}^{j\prime})$.

Finally, j's tentative payoff after any round is independent of the prices of items outside the set S_j. Hence, without affecting j's payoff, we may replace those other prices by prices $\mathbf{p}'$ so high that j would prefer not to acquire any of those other items at these prices. By mutual substitutes, $S_j \in X_j(p^j_{S_j}, p'_{L\setminus S_j})$, and hence $S_j \in D_j(p^j_{S_j}, p'_{L\setminus S_j})$. This implies that j's tentative profits are indeed nonnegative. Q.E.D.

The next issue is what happens when all bidders bid in a straightforward way. Theorem 2 provides an answer.

THEOREM 2. Assume that the licenses are mutual substitutes for all bidders and that all bidders bid straightforwardly. Then the auction ends with no new bids after a finite number of rounds. Let ($\mathbf{p}^*$, S^*) be the final standing high bids and license assignment. Then ($\mathbf{p}^*$, S^*) is a competitive equilibrium for the economy with modified valuation functions defined by $\hat{v}_j(T) = v_j(T) - \epsilon(\mathbf{p} \vee \hat{\mathbf{p}}) \cdot (TS^*_j)$ for each bidder j. The final assignment maximizes total value to within a single bid increment:

$$\max_{\{S_j\}} \sum_j v_i(S_j) - \sum_j v_j(S^*_j) \leq \epsilon(\mathbf{p}^* \vee \hat{\mathbf{p}}) \cdot L.$$

Proof. In view of theorem 1, at the end of every round, every bidder's tentative profit is positive. This implies that the total price of the licenses assigned to the bidders after any round of the auction is bounded above by the maximum total value of the licenses. Given the positive lower bounds on the bid increments, it follows that the auction ends after a finite number of rounds.

By construction, bidder j's demand at the final price vector $\mathbf{p}^*$ with j's modified valuation is the same as its demand at the corresponding personalized price vector $\mathbf{p}^j$ for the original valuation. Since there are no new bids by j at the final round, we may conclude from the condition of straightforward bidding and theorem 1 that $S^*_j \in \hat{D}(\mathbf{p}^*)$. Since this holds for all j, ($\mathbf{p}^*$, S^*) is a competitive equilibrium with the modified valuations.

For the second statement of the theorem, we can make the following calculation:

$$\begin{aligned}\max_S \sum_j v_j(S_j) &= \max_S \sum_j [\hat{v}_j(S_j) + \epsilon(\mathbf{p}^* \vee \hat{\mathbf{p}}) \cdot (S_j \backslash S_j^*)] \\ &\leq \max_S \sum_j [\hat{v}_j(S_j) + \epsilon(\mathbf{p}^* \vee \hat{\mathbf{p}}) \cdot S_j] \\ &= \max_S \sum_j \hat{v}_j(S_j) + \epsilon(\mathbf{p}^* \vee \hat{\mathbf{p}}) \cdot L \\ &= \sum_j \hat{v}_j(S_j^*) + \epsilon(\mathbf{p}^* \vee \hat{\mathbf{p}}) \cdot L \\ &= \sum_j v_j(S_j^*) + \epsilon(\mathbf{p}^* \vee \hat{\mathbf{p}}) \cdot L.\end{aligned}$$

The first equality follows from the definition of the modified valuations, the inequality from the restriction that all prices are nonnegative, and the following equality from the fact that S partitions L. The fourth step follows from the already proven fact that $(\mathbf{p}^*, S^*)$ is a competitive equilibrium for the modified valuations combined with the first welfare theorem and the fact that, with quasi-linear payoffs, a license assignment is efficient if and only if it maximizes the total value to all the bidders. Finally, the last equality follows by the definition of $\hat{v}_j(\cdot)$, which coincides with $v_j(\cdot)$ when evaluated at S_j^*. Q.E.D.

If the coefficient ϵ varies during the auction, then the most relevant values of ϵ for this analysis are ones that apply when bidders are last eligible to make new bids, which is normally near the end of the auction. (The activity rule is what makes this statement inexact.) This suggests that very high levels of *tatonnement* efficiency might be obtained by using small increments near the end of the auction. It was with this in mind that the Milgrom-Wilson rules originally adopted in the United States by the Federal Communications Commission (FCC) called for using smaller minimum bid increments in the final stage of the auction.[10]

The final questions in this section are, What relation does the auction outcome have to the competitive equilibrium outcome? Does a competitive equilibrium even exist in this setting with indivisible licenses? Theorem 3 provides answers.

THEOREM 3. Suppose that the licenses are mutual substitutes in

[10] That rule was later changed for reasons of transaction costs: smaller increments late in the auction led to large numbers of costly rounds with relatively little bidding activity.

demand for every bidder. Then a competitive equilibrium exists. For ϵ sufficiently small, the final license assignment $S^*(\epsilon)$ is a competitive equilibrium assignment.[11]

Proof. Let $\epsilon_n \rightarrow 0$ and let $S^*(\epsilon_n)$ and $\mathbf{p}^*(\epsilon_n)$ be corresponding sequences of final license assignments and prices. Since there are only finitely many possible license assignments, some assignment S^{**} must occur infinitely often along the sequence. Also, each license price is bounded above by the maximum value of a license package. So there exists a subsequence $n(k)$ along which $S^*(\epsilon_{n(k)}) = S^{**}$ and such that $\mathbf{p}^*(n(k))$ converges to some $\mathbf{p}^{**}$. By theorem 2, for all k, $S_j^{**} \in D_j(\mathbf{p}^*(n(k)), \epsilon_{n(k)})$, where the second argument of D_j identifies the relevant perturbed preferences. By the standard closed graph property of the demand correspondence, $S_j^{**} \in D_j(\mathbf{p}^{**})$, so $(S^{**}, \mathbf{p}^{**})$ is a competitive equilibrium. Q.E.D.

Intuitively, because the number of possible allocations is finite, a value-maximizing allocation generates a greater total value than the best nonmaximizing allocation by some amount $\delta > 0$. If the bid increment ϵ is sufficiently much smaller than δ, then, according to theorem 2, only an efficient allocation can result from straightforward bidding. The auction prices that support that allocation approximate the competitive equilibrium prices.

Thus, when all licenses are mutual substitutes for all bidders, the simultaneous ascending auction with straightforward bidding is an effective *tatonnement.* First, a bidder that bids straightforwardly during the auction is "safe": it is sure to acquire a set of licenses that is nearly optimal relative to its valuation and the final license prices, and it never risks actually losing money. If every bidder bids straightforwardly, then the auction eventually ends with an assignment that approximately maximizes the total value. Indeed, if the bid increment is small, then the final assignment exactly maximizes the total

[11] Milgrom and Roberts (1991) show the existence of a competitive equilibrium with mutual substitutes using a lattice-theoretic argument that does not require all goods to be divisible. They proceed to show that a wide variety of discrete and continuous "adaptive" and "sophisticated" price adjustment processes converge to the competitive equilibrium price vector. Unlike the present analysis, however, their analysis assumes that demand is given by a function, rather than by a correspondence, and they do not address the monotonicity of the auction process. Kelso and Crawford (1982) obtain results analogous to theorems 1 and 3 in a model of job matching. Gul and Stacchetti (1997) characterize utility functions that display "no complementarities," which is an alternative formulation of the idea of mutual substitutes. They also introduce a new auction process in which an auctioneer announces price vectors $\mathbf{p}$ and the bidders report their corresponding sets of demands $D_j(\mathbf{p})$. The auctioneer uses the reported information to control a continuous process of price increases. For the case of no complementarities in bidder utility, they demonstrate that their new auction process converges monotonically up to a competitive equilibrium.

TABLE 1

	A	B	AB
1	a	b	$a + b + c$
2	$a + d$	$b + d$	$a + b + d$

value and is a competitive equilibrium assignment. The final bids "approximately support" the solution, in the sense that they are close to the personalized prices that support the solution for each bidder. A number proportional to the bid increment bounds the error in each of these approximations.

The first three theorems were developed only for the case of licenses that are mutual substitutes. In practice, the status of spectrum licenses as substitutes or complements may often depend on how the licenses are defined. For example, in the DCS-1800 spectrum auction conducted in the Netherlands in February 1998, some of the offered licenses permitted use of only very small amounts of bandwidth relative to the efficient scale. A bidder that sought to establish an efficiently scaled mobile wireless telephone system would find that the value of, say, two small licenses is more than two times the value of a single license. Formally, that scale economy creates a complementarity among licenses: the value of a pair of licenses is more than the sum of the individual values. A similar complementarity from economies of scale and scope would be created by recent proposals in Australia to establish licenses covering small geographic areas with small amounts of bandwidth.

While there may be positive results available for some environments in which some of the goods are complements, my next result establishes a sharp limit. It shows that introducing into the previous model a single bidder for which licenses are not mutual substitutes leads to drastic changes in the conclusions.

THEOREM 4. Suppose that the set of possible individual valuation functions includes all the ones for which licenses are mutual substitutes in individual demand. Suppose that, in addition, the set includes at least one other valuation function. Then if there are at least three bidders, there is a profile of possible individual valuation functions such that no competitive equilibrium exists.

Intuition for theorem 4 is given in a two-license, two-bidder example, summarized in table 1. In the table, the licenses are denoted by A and B and the bidders by 1 and 2. Bidder 1 is the bidder for which licenses are not substitutes. This requires that the value of the pair AB exceed the sum of the individual values, that is, $c > 0$. Now let us introduce another bidder for which the same two licenses are substitutes. Let us take $c/2 < d < c$. In this case, the unique value-maximizing license allocation is for bidder 1 to acquire both licenses. In order to arrange for bidder 2 not to demand licenses, the prices must be $p_A \geq a + d$ and $p_B \geq b + d$, but at these prices bidder 1 is unwilling to buy the licenses. Consequently, there exist no equilibrium prices.

Proof of theorem 4. Suppose that there is a bidder in the auction with valuation function v for which licenses are not mutual substitutes. Then there is some price vector $\mathbf{p}$, real number $\epsilon > 0$, and licenses j and k such that $\{j, k\} \in X(\mathbf{p})$, but $j \notin X(\mathbf{p}\backslash(\mathbf{p}_j + \epsilon))$ and $k \notin X(\mathbf{p}\backslash(\mathbf{p}_j + \epsilon))$. For this bidder, define an indirect valuation function w on the set of licenses $\{j, k\}$ by

$$w(S) = \max_{T \subset L\backslash\{j,k\}} v(T \cup S) - \mathbf{p} \cdot T.$$

The bidder's demand for licenses in the set $\{j, k\}$ given the established prices $p_{L\backslash\{j,k\}}$ for the licenses besides j and k is determined by w. Set $a = w(j)$, $b = w(k)$, and $c = w(jk) - a - b$. From our assumptions about the bidder's demand, it follows that $c > 0$ and that $p_j + p_k < a + b + c < p_j + p_k + \epsilon$. Let us now introduce two new bidders whose values are given by the following valuation function:

$$\hat{v}(S) = \mathbf{p}(S\backslash\{j, k\}) + (a + d)\mathbf{1}_{j \in S} + (b + d)\mathbf{1}_{k \in S} - d\mathbf{1}_{j \in S, k \in S},$$

where $c/2 < d < c$. For the new bidders, the various licenses are mutual substitutes. (Indeed, the bidders' demands for each license in $L\backslash\{j, k\}$ are independent of all prices except the license's own price. For the two licenses j and k, the verification is routine.) By construction, the competitive equilibrium prices, if they exist, of licenses in $L\backslash\{j, k\}$ are given by p_{-jk}. But then the problem of finding market-clearing prices for j and k is reduced to the example analyzed above, in which nonexistence of equilibrium prices has already been established. Q.E.D.

This nonexistence is related as well to a problem sometimes called the "exposure problem" that is faced by participants in a simultaneous ascending auction. This refers to the phenomenon that a bidder that bids straightforwardly according to its demand schedule is exposed to the possibility that it may wind up winning a collection of licenses that it does not want at the prices it has bid, because the complementary licenses have become too expensive. If the bidders in the example in the table were to adopt only undominated strategies in every subgame of the simultaneous ascending auction game, then it is not possible that at the end of the auction bidder 1 will acquire both licenses unless the prices are at least $a + d$ and $b + d$ minus one increment. The reason is that bidder 2 always does at least as well (and could do better) in that subgame of the auction by placing one more bid. Whenever bidder 1 wins both licenses, it loses money, and at equilibrium it will anticipate that. To avoid "exposure" completely, bidder 1 must bid no more than a for license A and no more than b for license B. If it does so, then the outcome will be inefficient and the prices, a and b, will not reflect any of the potential "synergy" between the licenses.

One puzzle raised by the preceding analysis is that there have been spectrum auctions involving complements that appeared to function quite satisfactorily. The U.S. regional narrowband auction in 1994 was an auction in which several bidders successfully assembled collections of regional paging licenses in single spectrum bands to create the package needed for a nationwide paging service. In Mexico, the 1997 sale of licenses to manage point-to-point microwave transmissions in various geographic areas exhibited a similar pattern. What appears to be special about these auctions is that licenses covering different regions in the same spectrum band that were complementary for bidders planning nationwide paging or microwave transmission networks were not substitutes for any other bidders. The nonexistence theorem given above depended on the idea that licenses that are complements for one bidder are substitutes for another.[12]

The potential importance of the exposure problem is illustrated by the Netherlands auction mentioned earlier: the DCS-1800 auction in February 1998 in which 18 spectrum licenses were offered for sale. In that auction, two of the lots—designated A and B—were believed to be efficiently scaled; the remaining 16 lots were too small to be valuable alone and needed to be combined in groups of perhaps six licenses to be useful for a mobile telephone business. In

[12] The following table presents an example of nonexistence even when licenses are mutual complements for all bidders, but the degrees of complementarity vary. Tabulated are the values of three bidders (labeled 1, 2, and 3) for three licenses (A, B, and C). If a competitive equilibrium did exist, its assignment would be efficient, assigning licenses A and C to bidder 3 and B to bidder 1 or 2. For bidders 1 and 2 to demand their equilibrium assignments, the prices must satisfy $p_B \le 1$, $p_A + p_B \ge 3$, and $p_B + p_C \ge 3$. However, these together imply that $p_A + p_C \ge 4$, which is inconsistent with bidder 3's demand for the pair AC. So no competitive equilibrium exists.

	Licenses and Combinations						
Bidder	A	B	C	AB	AC	BC	ABC
1	1	1	0	3	1	1	3
2	0	1	1	1	1	3	3
3	1	0	1	1	3.5	1	3.5

this auction, the smaller licenses would naturally be complements for one another for bidders with no other licensed spectrum, but they would be substitutes for one another for bidders that were merely seeking to increase their amounts of licensed spectrum.

The outcome of the auction involved final prices per unit of bandwidth in lots A and B that were more than twice as high as those for any of the 16 smaller lots. It might seem that the bidders on lots A and B behaved foolishly since they might have acquired as much spectrum for less by bidding on smaller lots. However, bidders may have been deterred from aggressive bidding for the smaller lots for fear that that would drive up the prices of those lots while still leaving some of the winning bidders with too little bandwidth for an efficiently scaled business. This may simply be an instance in which, as suggested by theory, the prices fail to reflect the potential synergies among the licenses.

The problem of bidding for complements has inspired continuing research both to clarify the scope of the problem and to devise practical auction designs that overcome the exposure problem.

IV. Auctions and Game Theory

Another part of economic theory that has proved useful for evaluating alternative auction designs is game theory. Here I consider two such applications. The first model formalizes the ideas that motivated the introduction of the activity rule. The second is a study of how the auction closing rules affect the likelihood of collusive outcomes.

The Need for Activity Rules

In the design of the auction, one of the concerns was to estimate how long the auction would take to complete. This, in turn, depended on forecasting how aggressively bidders would behave. Could one count on the bidders to move the auction along, perhaps to economize on their own costs of participating? Or would the bidders sometimes have a strategic incentive to hold back, slowing the pace of the auction substantially?

There were several reasons to be skeptical that the bidders themselves could be relied on to enforce a quick pace. In the mutual substitutes model analyzed earlier, there is no affirmative gain to a bidder from bidding aggressively early in the auction, since all naive bidding paths lead to the same competitive equilibrium outcome. So bidders with a positive motive to delay might find little reason not to do so. In some of the spectrum auctions, the major bidders

TABLE 2

	License	
Bidder	A	B
1	15	30
2	10	Not eligible
3	Not eligible	5 with probability .9 15 with probability .1

included established competitors in the wireless industry that stood to profit from delays in new entry caused by delays in the auction process.

There can also be a variety of strategic motives for delay in the auction itself. Here I shall use a model to investigate one that is so common as to be decisive for planning the auction design. The model is based on the notion that the bidders are, or may be, budget-constrained.[13] (A large measure of strategic behavior in the actual spectrum auctions seemed to be motivated by this possibility.) If a bidder's competitor for a particular license is budget-constrained and its values or budget or both are private information, then the bidder may gain by concealing its ability or willingness to pay a high price until its competitor has already committed most of its budget to acquiring other licenses. The budget-constrained competitor may respond with its own delay, hoping to learn something about the prices of its highest-valued licenses before committing resources to other licenses. These behaviors delay the completion of the auction. What follows is a sample bidding game verifying that such behaviors are possible equilibrium phenomena.

Suppose that there are three bidders—1, 2, and 3—and two licenses—A and B. Each bidder has a total budget of 20, and its total payments cannot exceed this limit. A bidder's payoff is its value for the licenses it acquires minus the total amount it pays. The values of the three bidders for the two licenses are listed in table 2.

The rules of the game are as follows. Initially, the prices are zero, and both items are assigned to the auctioneer. At any round, a bidder can raise the bid by one unit on any license for which it is eligible to bid. Ties are broken at random. After a round with no new bids, the auction ends. Payoffs are determined as described above.

[13] Budget constraints can have profound effects on bidding behavior and equilibrium strategies. Pitchik and Schotter (1988) initiated research into the effects of budget constraints; see also Che and Gale (1996, 1998). For some of the other effects of budget constraints on actual bidder behavior in the spectrum auctions, see Milgrom (1995, chap. 1).

My question is, Does there exist a (sequential) equilibrium in which bidders 2 and 3 bid "straightforwardly," that is, in which each raises the bid on a license whenever it is not assigned the license and its value strictly exceeds the current highest bid? If bidder 3's value is common knowledge among the bidders, then one can routinely verify that the answer is affirmative. Bidder 1's corresponding strategy depends on bidder 3's value for license B. If that value is 5, then at the equilibrium, bidder 1 bids in the same straightforward manner as the other two bidders. If, however, bidder 3's value is 15, then bidder 1's best reply is different. At one equilibrium, 1 bids straightforwardly on license B and limits its bids on license A to ensure that it will win license B with its limited budget.

If 3's value is private information, however, then the answer changes. For suppose that bidders 2 and 3 bid straightforwardly. Then 1 could learn 3's value by bidding on license B until it was assured of acquiring that license, then devoting its remaining budget in an attempt to win license A. In particular, 1 would always win license B. It would also win license A at a price of 10 or 11 when 3's value for B was low. There can be no equilibrium with these properties, however. For if there were, then when bidder 3 has the high value, it could wait until 1 bids 10 or 11 on license A before bidding more than 5 on license B. Then 3 would win license B and earn a positive profit.

THEOREM 5. There is no sequential equilibrium of the private information game in table 2 in which bidders 2 and 3 each bid "straightforwardly," as described above.

Both bidders 2 and 3 may have an incentive to slow their bidding in this auction, each hoping that bidder 1 will become unable to compete effectively for one license because it has spent its budget on another license. What the equilibrium in this example does not show is a delay induced by bidder 1, since it avoids committing resources until after bidder 3 has shown its hand. I conjecture that the example can be extended to incorporate that feature, so that all bidders have a tendency to delay.

In the actual spectrum auctions, the activity rule limited such wait-and-see strategies by specifying that a bidder that remained inactive in the early rounds of the auction would be ineligible to bid in later rounds. However, the first auctions cast doubt on the necessity of the rule. In the national and regional narrowband auctions, there was far more bidding activity than required by the activity rule, leading some to propose that the auction be simplified by dropping the rule. However, the AB block PCS auction, which was the third simultaneous ascending auction, followed quite a different pattern.

For the AB auction, the volume of activity associated with each

TABLE 3

	A	B	AB	Budget
1	4	. . .	. . .	3
2	. . .	4	. . .	3
3	$1 + \epsilon$	$1 + \epsilon$	$2 + \epsilon$	2

license is measured by the population in the region covered by the licenses according to the 1990 U.S. census ("POPs"). The average license in this auction covered a region with approximately 5 million POPs. The auction generated 3,333 data points, each consisting of a vector of bids made by a bidder at a round.[14] Only 30 of the 3,333 observations reveal activity that exceeds the required level by at least one average size license, that is, 5 million POPs, and only 140 observations reveal activity exceeding required activity by more than 1 million POPs. Thus bidders most often bid only slightly more than was minimally necessary to maintain their current bidding eligibility.

Free Riding

One of the main issues in the early debates about the spectrum auction was whether all bidding should apply to individual licenses or whether, instead, bids for combinations of licenses should be allowed. According to one combinatorial bidding proposal, bids would first be accepted for certain predetermined packages of licenses, such as a nationwide collection of licenses, and then bidding on individual licenses would ensue. After all bidding had ceased, the collection of bids that maximize total revenues would be the winning bids, and licenses would be assigned accordingly.[15] The model of this auction below assumes that in the event of ties, package bids are selected in preference to bids on individual licenses and that bids must be entered as whole numbers.

The primary economic argument against allowing combination bids is that such bids can give rise to a free-rider problem among bidders on the individual licenses, leading to avoidable inefficiencies. Table 3 provides a simplified version of an example I presented

[14] Observations in which bidders take a "waiver" are excluded for two reasons. First, the required activity does not apply at rounds with waivers, so there is no natural x variable. Second, each bidder that ceases bidding before the end of the auction automatically exercises five waivers according to the FCC rules, so those observations contain no information about bidder decision making.

[15] Depending on what combinations are allowed, there may also need to be rules specifying the winner when there are overlapping combinations. Generally, the recommendation was that the winning set of bids should be the set that maximizes the total bid price.

during the deliberations to show how that can happen. In this example, there are three bidders—labeled 1, 2, and 3—and two licenses—A and B. Bidders 1 and 2 are willing to pay up to 4 for licenses A and B, respectively, and neither is eligible to acquire the other license.[16] With ϵ small and positive, bidder 3 has the lowest values for the licenses but is distinguished by its desire to acquire both. To keep the strategy spaces small and ease the analysis, I impose economically insignificant budget constraints on the bidders, as shown in table 3.

With the specified values, the sole efficient license assignment has bidders 1 and 2 acquiring licenses A and B, respectively. With bids restricted to be whole numbers, that corresponds to a subgame-perfect equilibrium of the simultaneous ascending auction. At the equilibrium, bidders 1 and 2 make minimum bids at each round as necessary to acquire their respective licenses of interest, whereas bidder 3 bids 1 for each license and then gives up.

If the proposed combinatorial auction is used, bidder 3 can refrain from bidding for licenses A and B directly, bidding instead for the pair AB. This strategy creates a free-rider problem for bidders 1 and 2. A high bid by bidder 1 on license A helps bidder 2 to acquire license B. A symmetric observation applies to bidder 2. Each would prefer that the other raise the total of the individual bids sufficiently to beat 3's bid.

Even in the complete information case shown here, this free-rider problem can lead to inefficient mixed-strategy equilibria. The corresponding equilibrium strategies are as follows. In the combination bidding round, bidder 3 bids 2 for the license combination AB. Bidder 1 raises the price of license A by 1 whenever it does not own the standing high bid for that license. Otherwise, if at any time during the auction the license prices are 1 for A, 1 for B, and 2 for AB, then bidder 1 raises its high bid on license A with probability two-thirds. Bidder 2's strategy is symmetrical to bidder 1's but is focused on license B instead of license A.

The key to understanding this equilibrium is to recognize the payoffs in the subgame after the prices are 1 for A, 1 for B, and 2 for package AB. The payoff matrix for bidders 1 and 2 in that subgame is shown in table 4.

This subgame has a symmetric equilibrium in which each bidder raises the bid with probability two-thirds. Backward induction from

[16] In the actual auctions, bidders were ineligible to acquire additional wireless telephone licenses for areas they already served. This restriction was motivated by competition policy.

TABLE 4

	Raise Bid	Don't Raise
Raise bid	2, 3	2, 3
Don't raise	3, 2	0, 0

there supports the equilibrium strategies described above. At the equilibrium, there is a one-ninth probability that 3 acquires both licenses, even though its value for those licenses is just one-fourth of the total of the competitors' values. This example is representative of a robust set of examples, including especially ones with asymmetric information that make the free-rider problem even harder to resolve.

The following theorem summarizes this discussion.

THEOREM 6. The proposed two-stage auction (in which combinatorial bidding is followed by a simultaneous ascending auction for individual licenses) can introduce inefficient equilibrium outcomes that would be avoided in the simultaneous ascending auction without combinatorial bidding.

It bears emphasis that this defect applies to the particular combinatorial rule that was proposed and is not a general criticism of all combinatorial bidding.

Collusion and Closing Rules

Motivated by the idea of the *tatonnement,* the rules of the spectrum auction specified that bidding would close on no licenses until there were no new bids on any license. In that way, if a license that changed hands at some round were a substitute or complement for another license, the losing bidder could react by bidding for the substitute or withdrawing a bid for a complement, and the winner could react in the reverse way.

Strategically, however, simultaneous closings create opportunities for collusion that can be mitigated by other closing rules.[17] To illustrate this in a simple model, suppose that there are two bidders, 1 and 2, and two licenses, A and B. Each bidder has a value for each license of 10. The auction rules are the same as in the preceding

[17] An unpublished paper by Rob Gertner (1995) inspired my analysis of closing rules. His presentation analyzed the vulnerability to collusion of the simultaneous ascending auction with simultaneous closings and showed that the same form of collusion is not consistent with equilibrium in the traditional auctions in which items are sold one at a time, in sequence.

subsection, with a simultaneous close of bidding on all licenses when there is no bidding on any license. The next two theorems, the proofs of which are straightforward, show that both "competitive" and "collusive" outcomes are consistent with equilibrium in this game.

THEOREM 7. The following strategy, adopted by both bidders, constitutes a sequential equilibrium of the game with simultaneous closes of bidding: if the price of either license is below 10, bid again on that license.

This is the "competitive" outcome and results in prices of 10 for both licenses and zero profits for the bidders. However, other outcomes are also possible.

THEOREM 8. The following strategies constitute a sequential equilibrium of the game with simultaneous closes of bidding. (1) For bidder 1, if 2 has never bid on license A, then if license A has received no bids, bid $1 on license A; otherwise, do not bid. If 2 has ever bid on license A, then bid according to the strategy described in theorem 7. (2) Bidder 2 bids symmetrically.

This is the most collusive equilibrium, resulting in prices of just 1 for each license and total profits of 18 for the two bidders, which are the lowest prices possible if the licenses are to be sold. The collusive outcome is supported by the threat, inherent in the strategies, to shift to competitive behavior if the other party to the arrangement does not refrain from bidding on a particular license.

An extreme alternative is to close bidding on a license after any round in which there is no new bid on *that license*. This rule excludes the possibility that bidders can each retaliate if the other cheats on the arrangement. For example, suppose that the auction is supposed to end after round n with a bid price of $b \leq 8$ on license A, won by bidder 1. Then bidder 2 has nothing to lose and, in the trembling-hand logic of equilibrium, something to gain by raising the price at round $n + 1$. Consequently, we have the following result.

THEOREM 9. In the game with license-by-license closes of bidding, at every (trembling-hand) perfect equilibrium, the price of each license is at least 9.

Similar results can be obtained from a rule that arranges for bidding to close on a license if there has been no new bid in the past three rounds. Alternatively, bidding may close on a license when there has been no new bid for three rounds and the total number of new bids on all licenses for the past five rounds is less than some trigger value. Rules along these lines can allow for substitution among licenses until late in the auction while still deterring some of the most obvious opportunities for collusion.

V. Dynamic Bidding for Combinations of Licenses

The considerations raised in the *tatonnement* analysis suggest the need to use a mechanism that does not rely simply on prices for individual licenses and instead allows bidding for license packages. An auction design that, in theory, uses combination bidding to good effect is the generalized Vickrey auction, also called the Groves-Clarke "pivot mechanism" (Vickrey 1961; Clarke 1971; Groves and Loeb 1975). Since that will serve as our standard of comparison, I review it briefly here.

Let L denote the set of available licenses and let P be the set of license assignments; these are indexed partitions of L. For any assignment $S \in P$, partition element S_i represents the set of licenses assigned to bidder i.

The rules of the generalized Vickrey auction are as follows. Each bidder submits a bid that specifies a value for every nonempty subset of L. For any set of licenses T, let $v_i(T)$ denote i's bid for that set. The auctioneer chooses the license assignment S^* that maximizes $v_1(S_1^*) + \ldots + v_N(S_N^*)$. Each bidder i pays a price p_i for its licenses according to the formula

$$p_i = \max_{S \in P} \sum_{j \neq i} v_j(S_j) - \sum_{j \neq i} v_j(S_j^*).$$

It is well known that, subject to certain assumptions,[18] the bidders in a generalized Vickrey auction have a dominant strategy, which is to set their bids for each license package equal to its actual value. When each bidder uses its dominant strategy, licenses are assigned efficiently. Moreover, if the bidder types have independent, atomless[19] distributions, then any other auction design that leads to effi-

[18] Among the important assumptions are the following. First, the bidders know their own values; i.e., this is a pure private-value model with no common-value elements. (See Milgrom and Weber [1982] for a discussion of this assumption.) Second, bidders must care only about the sets of licenses they acquire and the prices they pay, and not about the identities of the other license acquirers and the prices they pay (although extensions of the Vickrey auction can accommodate bidders that care about the entire license allocation). Third, budget constraints must never be binding. Each of these assumptions is a strong one. None precisely fits the facts about the U.S. spectrum auctions. In addition, there is the relatively more innocuous assumption that bidder preferences are quasi-linear. This means that a bidder's utility is representable as the value of the licenses assigned to it minus the price that it pays.

[19] I am indebted to Paul Klemperer for pointing out the necessity of the atomless type distribution condition. In this application, a "type" is a vector of values for licenses and combinations of licenses.

cient outcomes must involve the same expected payments by all the types of all the bidders (see, e.g., Engelbrecht-Wiggans 1988).

The generalized Vickrey auction itself is not practical for use in spectrum sales. If there were no restrictions on feasible license combinations, the number of combinations would be $2^{|L|} - 1$. Most of the sales being conducted presently involve hundreds of licenses, and even though in practice most of the combinations can be ruled out as infeasible or irrelevant, the number of potentially important combinations is still infeasibly large.[20] I seek to use the Vickrey auction here as a benchmark, in much the same way that the competitive equilibrium benchmark is used in market welfare analyses.

Given that it is infeasible to specify all relevant combinations in advance, one idea to economize on computing power is to specify combinations as the auction progresses. The leading such proposal is based on a procedure called the "adaptive user selection mechanism" (AUSM) that was developed in experimental economics laboratories for solving what the experimenters regarded as "difficult" resource allocation problems (Banks, Ledyard, and Porter 1989; Ledyard, Noussair, and Porter 1996).

The AUSM differs from the simultaneous ascending auction in a number of respects, and many of its features have been proposed for adoption in the spectrum auctions. Among the proposed changes are the following: First, allow bidding to take place continuously in time, rather than force bidders to bid simultaneously in discrete rounds. Second, in place of an activity rule, follow the experimenters' technique of using random closing times, which motivate bidders to be active before the end of the auction. Third, permit bids for combinations of licenses rather than just for individual licenses. When a new combinatorial bid is accepted, it displaces all previous standing high bids for individual licenses or combinations of licenses that overlap the licenses in the new bid. The new bid should be accepted if the amount of the bid is greater than the sum of the displaced bids. Fourth, allow the use of a "standby queue" on which bidders may post bids that cannot, by themselves, displace existing bids but become available for use in new combinations. For example, suppose that bidder 1 owns the standing high bid of 20 for license combination ABCD. Bidder 2 is interested in acquiring AB for a price of up to 15 but has no interest in CD. It may post a bid of 12

[20] An additional objection to Vickrey auctions is that they require bidders to reveal their value estimates. Bidders have been reluctant to do that, possibly because they fear that reporting their values would reveal information to competitors about how they form estimates, what discount rates they use, what financing they have available, or what their business plans are.

TABLE 5

Bidder	License ID A	B	AB	Vickrey Price
1	V_1	V_1	V_1	$V_3 - V_2$
2	V_2	V_2	V_2	$V_3 - V_1$
3	0	0	V_3	$V_1 + V_2$

for AB on the standby queue. Suppose that it does so and that bidder 3 is willing to pay up to 15 for CD. Then bidder 3 may "lift" 2's bid from the standby queue and submit that together with its bid of 10 for license combination CD, thereby creating a bid of 22 for the combination ABCD. Under the rules, bidders 2 and 3 become the new owners of the standing high bids.

We can begin to analyze this proposal using a simple example, represented in table 5. There are three bidders, labeled 1–3, and two licenses. The first two bidders each want to acquire a single license; the third bidder is interested only in the pair. The final column shows what price the bidder would pay in a generalized Vickrey auction in which it is a license winner.

The bidders' values are drawn from continuous distributions. For the first two bidders, the distribution has support on $[a, b]$, and for the third bidder, it has support on $[c, d]$. We assume that $2a < d$ and that $2b > c \geq b$. These inequalities mean that (1) there is a priori uncertainty about the efficient license assignment, and (2) the two single-license bidders need to coordinate to be able to outbid bidder 3.

Since there are many different implementations of AUSM, I regard it as a class of games. I limit attention to implementations in which bidding takes place in rounds and does not end after a round in which there are new bids. I look for properties of equilibrium in undominated strategies of any such AUSM game in which no bidder makes jump bids. Three general properties hold. First, no bidder j bids more than its own actual value V_j, for to do so would entail using a weakly dominated strategy. With no jump bids, this implies that bidder 3 never pays more than $V_1 + V_2$. Second, since bidder 3 always has an opportunity to respond to the bids by 1 and 2, equilibrium entails that bidder 3 wins a license when $V_3 > V_1 + V_2$. Free-riding among the individual bidders may mean that bidder 3's AUSM equilibrium price is strictly less than the Vickrey price $V_1 + V_2$. Third, when the single-license bidders 1 and 2 win licenses in an AUSM game, the total price they pay is V_3. They win only when $V_1 + V_2 >$

V_3, and given the free-rider problem, they may not always win even when that inequality holds.[21] From the preceding inequality, the total price V_3 that the bidders pay when they win is strictly greater than the total Vickrey price of $2V_3 - V_1 - V_2$. This leads to the following conclusion.

THEOREM 10. In the example analyzed here, the total equilibrium prices under AUSM for the single-license bidders are always at least as high as and sometimes higher than the Vickrey prices, whereas the price paid by the combination bidder is never more and sometimes less than the Vickrey price. The combination bidder wins (weakly) more often than it would at an efficient auction, and the single-license bidders win (weakly) less often than they would at such an auction.

Experiments have established that AUSM performs well in some environments with significant complementarities. The questions for auction designers are, Which kinds? And how can their disadvantages be minimized? Identifying biases is a first step toward answering such questions.

VI. Two Additional Questions

One of the most frequently expressed doubts about the spectrum auctions is the doubt that the form of the auction matters at all. After all, the argument goes, one should expect that if the initial assignment resulting from the auction is inefficient and if licenses are tradable, the license owners will be motivated after the auction to buy, sell, and swap licenses until an efficient assignment is achieved.

There are both theoretical and empirical grounds for rejecting this argument. The theoretical argument is developed at length in Milgrom (1995). Briefly, the argument combines two theoretical observations from the theory of resource allocation under incomplete information in private-values environments. The first observation is that, once property rights have been assigned, ex post bargaining cannot generally achieve efficient rearrangement of the rights. The older theoretical literature shows this for the case in which there are just two parties to the bargain and the efficient allocation of the license is uncertain. Recent work by Cai (1997) suggests that the efficient outcomes become even less likely when there are multiple parties involved, as is the case when a bidder needs to assemble a

[21] Notice that a solution to the free-rider problem may require that one bidder pay more for its license than another bidder pays for a perfectly substitutable license. One may guess that such a solution would be particularly difficult to achieve if the bidders are ex ante identically situated.

collection of spectrum licenses from multiple owners to offer the most valuable mobile telephone service.[22] The years of delay in developing nationwide mobile telephone services in the United States, despite the value that customers reportedly assign to the ability to "roam" widely with their phones, testify to the practical importance of this theoretical effect. An inefficient initial assignment cannot, in general, be quickly corrected by trading in licenses after the auction is complete.

In contrast, the generalized Vickrey auction applied to the initial assignment of rights in the same environment can achieve an efficient license assignment—at least in theory. There are significant practical difficulties in implementing a Vickrey auction in the spectrum sales environment, but the theoretical possibility of an auction that always yields an efficient assignment establishes the possibility that a good initial design can accomplish objectives that ex post bargaining cannot.

A second common question concerns the trade-off between the goals of allocational efficiency and revenue. The primary goal of the spectrum auctions was set by the 1993 budget legislation as one of promoting the "efficient and intensive use" of the radio spectrum. However, the simultaneous ascending auction is now also being touted for other applications, such as the sale of stranded utility assets (Cameron, Cramton, and Wilson 1997) in which revenue is regarded as an important objective. Such applications call for putting more emphasis both on how the auction rules affect revenue and on the extent of the conflict between the goals of efficiency and revenue in multiobject auctions.

Particularly when the number of bidders is small, the goals of efficiency and revenue can come into substantial conflict. A particularly crisp example of this is found in the decision about how to package groups of objects when there are only two bidders.[23] Using the spectrum sale as an example, suppose that the available bands of spectrum are denoted $\{1, \ldots, B\}$ and that these are packaged in licenses $\mathbf{L} = \{1, \ldots, L\}$. The jth license consists of a set of bands $S_j \subseteq \{1, \ldots, B\}$, and a "band plan" is a partition $S = \{S_1, \ldots, S_N\}$ of the L bands into $N \leq L$ licenses.

Next, I introduce a special assumption. Suppose that each bidder i's valuation for any license is given by $X_i(S_j) = \sum_{k \in S_j} x_{ik}$. This assump-

[22] The same theoretical analysis applies to attempts to resolve the problem by contracting: ex post bargaining under incomplete information after property rights have already been assigned does not generally lead to efficient outcomes.

[23] See Palfrey (1983) for a related analysis, showing that bundling can increase revenue even when it reduces efficiency in various kinds of auctions.

tion abstracts from some potential interactions between efficiency and revenue and isolates the one effect on which I wish to focus.

Let $R(S)$ denote the revenue from the license sales corresponding to the band plan S, and let $V(S)$ be the total value of the licenses to the winning bidders when the licenses are sold individually in simultaneous second-price auctions and each bidder adopts its dominant strategy.

The conflict between efficiency and revenue in this context is very sharp. When one is choosing band plans in this setting, there is a *dollar-for-dollar trade-off* between the seller's revenue $R(S)$ and the value $V(S)$ of the final license assignment: any change in the band plan S that increases the value of the assignment reduces the seller's revenue by an equal amount!

THEOREM 11. The sum of the value created and the revenue generated by the auction is a constant, independent of the band plan S: $R(S) + V(S) = X_1(\mathbf{L}) + X_2(\mathbf{L})$. Coarser band plans generate higher revenues and create less value.

Proof. For the first statement, it suffices to show that, for any license S_j, the value created by the auction plus the license price is equal to $X_1(S_j) + X_2(S_j)$, for the result then follows by summing over licenses.

Suppose (without loss of generality) that bidder 1 has the higher value for the license. Then in an English auction, bidder 1 will win; the winner's value will be $X_1(S_j)$; and the price will be the second-highest value, $X_2(S_j)$.

For the second statement, recall that the outcome of the ascending auction is to assign each license to the bidder that values it most highly. Given two band plans S and S', with S coarser than S', the associated values are

$$
\begin{aligned}
V(S) &= \sum_{T\in S} \max\left(\sum_{k\in T} x_{1k}, \sum_{k\in T} x_{2k}\right) \\
&\leq \sum_{T\in S} \sum_{\substack{T\subseteq T \\ T\in S}} \max\left(\sum_{k\in T} x_{1k}, \sum_{k\in T} x_{2k}\right) = V(S).
\end{aligned}
$$

The inequality applies term by term to the maxima over sets $T \in S$. Q.E.D.

To illustrate the theorem, suppose that there are two bands with $x_{11} > x_{21}$ but $x_{12} < x_{22}$, and suppose in addition that $x_{11} + x_{12} > x_{21} + x_{22}$. There are two possible band plans according to whether the bands are sold as one license or two. When the bands are sold separately, bidder 1 wins band 1 at price x_{21} and bidder 2 wins band 2 at price x_{12}, creating a total value of $x_{11} + x_{22}$ and revenue of $x_{21} + x_{12}$. When the bands are sold together, bidder 1 acquires both at price $x_{21} + x_{22}$, creating a total value of $x_{11} + x_{12}$. The loss of value from adopting this plan is $x_{22} - x_{12}$, which is precisely the same as the increase in revenue from the same change.

In the analysis of Cameron et al. (1997), the items being sold are electrical generating plants or other "stranded utility assets" associated with deregulation. In that case, revenue (which reduces the burden on ratepayers) and efficiency are both typically among the goals of the public authority. In that case, if the number of serious bidders is sufficiently small, then the effect identified in this suggestion contributes to a trade-off in the public decision process between the goals of revenue and efficiency.

VII. Conclusion

In the last few years, theoretical analyses have clearly proved their worth in the practical business of auction design. Drawing on both traditional and new elements of auction theory, theorists have been able to analyze proposed designs, detect biases, predict shortcomings, identify trade-offs, and recommend solutions.

It is equally clear that designing real auctions raises important practical questions for which current theory offers no answers. The "bounded rationality" constraints that limit the effectiveness of the generalized Vickrey auction are important ones and have so far proved particularly resistant to simple analysis. Because of such limits to our knowledge, auction design is a kind of engineering activity. It entails practical judgments, guided by theory and all available evidence, but it also uses ad hoc methods to resolve issues about which theory is silent. As with other engineering activities, the practical difficulties of designing effective, real auctions themselves inspire new theoretical analyses, which appears to be leading to new, more efficient and more robust designs.

References

Ausubel, Lawrence M., and Cramton, Peter. "Demand Reduction and Inefficiency in Multi-unit Auctions." Working Paper no. 96-07. College Park: Univ. Maryland, 1996.

Banks, Jeffrey S.; Ledyard, John O.; and Porter, David P. "Allocating Uncertain and Unresponsive Resources: An Experimental Approach." *Rand J. Econ.* 20 (Spring 1989): 1–25.

Cai, Hong-Bin. "Delay in Multilateral Bargaining under Complete Information." Working paper. Los Angeles: Univ. California, 1997.

Cameron, Lisa; Cramton, Peter; and Wilson, Robert. "Using Auctions to Divest Generating Assets." *Electricity J.* 10, no. 10 (1997): 22–31.

Che, Yeon-Koo, and Gale, Ian. "Financial Constraints in Auctions: Effects and Antidotes." In *Advances in Applied Microeconomics*, vol. 6, edited by Michael R. Baye. Greenwich, Conn.: JAI, 1996.

———. "Standard Auctions with Financially Constrained Bidders." *Rev. Econ. Studies* 65 (January 1998): 1–21.

Clarke, Edward H. "Multipart Pricing of Public Goods." *Public Choice* 11 (Fall 1971): 17–33.

Cramton, Peter C. "Money Out of Thin Air: The Nationwide Narrowband PCS Auction." *J. Econ. and Management Strategy* 4 (Summer 1995): 267–343.

———. "The FCC Spectrum Auctions: An Early Assessment." *J. Econ. and Management Strategy* 6 (Fall 1997): 431–95.

Engelbrecht-Wiggans, Richard. "Revenue Equivalence in Multi-object Auctions." *Econ. Letters* 26, no. 1 (1988): 15–19.

Gertner, Robert. "Revenue and Efficiency Differences between Sequential and Simultaneous Multiple-Unit Auctions with Limited Competition." Manuscript. Chicago: Univ. Chicago, Grad. School Bus., 1995.

Groves, Theodore, and Loeb, Martin. "Incentives and Public Inputs." *J. Public Econ.* 4 (August 1975): 211–26.

Gul, Faruk, and Stacchetti, Ennio. "Walrasian Equilibrium without Complementarities." Working Paper no. 97-01. Ann Arbor: Univ. Michigan, Dept. Econ., 1997.

Kelso, Alexander S., Jr., and Crawford, Vincent P. "Job Matching, Coalition Formation and Gross Substitutes." *Econometrica* 50 (November 1982): 1483–1504.

Ledyard, John O.; Noussair, Charles; and Porter, David. "The Allocation of Shared Resources within an Organization." *Econ. Design* 2 (November 1996): 163–92.

McAfee, R. Preston, and McMillan, John. "Analyzing the Airwaves Auction." *J. Econ. Perspectives* 10 (Winter 1996): 159–75.

Milgrom, Paul. "Auction Theory for Privatization." Book manuscript. Stanford, Calif.: Stanford Univ., 1995.

Milgrom, Paul, and Roberts, John. "Adaptive and Sophisticated Learning in Normal Form Games." *Games and Econ. Behavior* 3 (February 1991): 82–100.

Milgrom, Paul, and Weber, Robert J. "A Theory of Auctions and Competitive Bidding." *Econometrica* 50 (September 1982): 1089–1122.

Palfrey, Thomas R. "Bundling Decisions by a Multiproduct Monopolist with Incomplete Information." *Econometrica* 51 (March 1983): 463–83.

Pitchik, Carolyn, and Schotter, Andrew. "Perfect Equilibria in Budget-Constrained Sequential Auctions: An Experimental Study." *Rand J. Econ.* 19 (Autumn 1988): 363–88.

Vickrey, William. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *J. Finance* 16 (March 1961): 8–37.