

PACKAGE AUCTIONS AND EXCHANGES

BY PAUL MILGROM¹

We report recent advances concerning the package allocation problem, in which traders seek to buy or sell combinations of goods. The problems are most difficult when some goods are not substitutes. In that case, competitive equilibrium typically fail to exist but the core is non-empty and comprises the competitive solutions. Also in that case, the Vickrey auction fails to select core allocations and yield revenues that are less than competitive. The Ausubel-Milgrom auction generally selects core allocations and, when goods are substitutes, prescribes the Vickrey allocation. We also evaluate the problems and promise of mechanisms for the package exchange problem.

KEYWORDS: Combinatorial bidding, package auction, stable match, core.

1. INTRODUCTION

THIS LECTURE REPORTS on the state of auction and exchange design for a class of challenging resource allocation problems—*package allocation problems*—that are common in both the private and public sectors of the economy. The class includes problems in which the goods to be traded are discrete and for which some agents might wish to acquire several items. It includes both problems in which the bidders are buyers and the seller may have some flexibility about what to sell, and ones in which the bidders are sellers whose bids may specify not merely prices for various packages, but also the attributes of the goods offered. Packages and attributes can interact in complicated ways in the buyer's valuations. In short, this is about resource allocation without some of the simplifying assumptions that are traditionally made for neoclassical economic analysis.

Most of the known results, and most of my attention in this paper, focus on the special case in which one party—a buyer or a seller—is passive and receives the residual. This is the *auction design problem*. Economists working in the tradition of Walras have often approached auctions differently, as a means for discovering market-clearing prices, but that approach is not especially helpful for the problems studied here. For the package allocation problem, market-clearing prices cannot be guaranteed to exist except when goods are substitutes for all participants. The nonexistence of market-clearing prices poses a definitional problem for public sector sales in which the government seeks to buy (or sell) goods at “competitive” terms. I subsequently argue that the *core* allocations, which include any competitive equilibrium allocations when such exist, coincide exactly with the competitive equilibrium allocations in a reformulated

¹This article is based on the Fisher–Schultz Lecture presented at the 2004 Econometric Society European Meeting. This paper has benefited from the suggestions of many people. Special thanks to Dan Quint, John Hegeman, the editor, and the referees for their detailed comments. Financial support for this research was provided by the National Science Foundation under grant ITR-0427770.

version of the problem where all participants' resources are sold as packages, and so provide a useful standard for evaluating package allocation mechanisms for public sector applications.

After examining when competitive equilibrium exists and how the core characterizes competition, I investigate the performance of the Vickrey auction. The well-known appeal of the Vickrey auction is that it is the only incentive-compatible mechanism that always leads to efficient outcomes on the class of transferable-utility environments. In practice, however, Vickrey auctions are never used to solve package allocation problems and that would seem to be a puzzle for economic theory. I explore several potential explanations that have been offered and conclude that none is convincing, because other mechanisms with the same defects are used in practice. I argue that a critical disadvantage of the Vickrey auction is that when goods are not substitutes for the bidders, the auction can lead to uncompetitively low prices for a seller (more generally, low payoffs for the auctioneer) even when all the bids are high. This happens because revenues in a Vickrey auction fail to be monotonically nondecreasing in the bids, and that nonmonotonicity also exposes the auction to several other kinds of manipulations. Sellers in a Vickrey auction can sometimes increase revenues by disqualifying bidders; buyers can sometimes reduce prices and enhance their profits by sponsoring fake bidders; and joint deviations by *losing* bidders who raise their bids can allow them to become winners at very low prices. These possibilities create threats that further diminish the usefulness of the Vickrey mechanism.

These failings of the Vickrey auction can be ruled out in environments in which goods are substitutes for all the bidders, but not for any broader class of environments. Notice that the substitutes condition is the same one that guarantees the existence of market-clearing prices. Because the absence of market-clearing prices is perhaps the foremost reason to adopt a complicated package mechanism (rather than an auction that prices goods individually), the conclusion that the failings of Vickrey design emerge in precisely those environments is discouraging for applications.

How do these failings of the Vickrey auction compare with those of other package-auction mechanisms? The mechanism most commonly used in actual practice is the sealed-bid package auction in which the winning bidders pay the amounts of their winning bids. Although this mechanism is not incentive-compatible, it also suffers none of the latter mechanism's major failings. In the pay-as-bid mechanism, higher bids lead to higher prices, sellers cannot benefit by disqualifying bidders, it is impossible to increase profits by sponsoring shells, and a coalition of losing bidders can become winners only by paying more than the current winning bidders.

The pay-as-bid package-auction mechanism is identical to the *menu auction* analyzed by Bernheim and Whinston (1986), who studied the full-information

equilibrium outcomes.² They identified equilibrium outcomes with the *bidder-optimal* allocations, which are core allocations with the property that there exists no other core allocation that raises some bidder's payoff without reducing another bidder's payoff. This performance provides a partial standard against which to test other mechanisms. However, because this analysis relies on complete information, it is not clear that anything helpful can be inferred from it about performance in the incomplete-information cases that dominate modern mechanism design analysis.

A third proposal for a package-auction mechanism is the *ascending proxy auction* of Ausubel and Milgrom (2002). This mechanism has several appealing theoretical properties. First, when goods are substitutes, truthful reporting of values is an ex post equilibrium of the mechanism, which duplicates the performance of the Vickrey auction in precisely the class of cases where the Vickrey mechanism is guaranteed to lead to core outcomes. Second, it has the property that it selects core allocations with respect to reported preferences, so it avoids the low revenue problems of the Vickrey mechanism. Finally, its full-information equilibrium outcomes and strategy profiles include the ones that Bernheim and Whinston select for their menu auction. In summary, the ascending proxy auction mechanism combines the known theoretical advantages of the two most studied package-auction mechanisms.

After treating the package-auction problem, we move on to the harder problem of the package exchange. In a famous contribution, Myerson and Satterthwaite (1983) showed that even for a simple bilateral exchange, there is no incentive-compatible mechanism that always implements an efficient allocation with balanced budgets. In that problem, however, the full-information core is always nonempty. In contrast, in a package exchange problem with multiple parties on both sides, the core can be empty.³ With an empty core, voluntary participation, and fully informed participants, there always exists some coalition of buyers and sellers that can do better on its own than by playing in whatever mechanism has been proposed. These are daunting problems for the design of practical exchanges and suggest that one should not expect too much from package exchange designs.

The ascending proxy auction can be adapted to the problem of package exchanges provided the regulator who operates the mechanism is willing to accept cash payments as part of the outcome. There are, however, no positive incentive results for such a mechanism. An alternative approach that holds much promise has been devised by Parkes, Kalagnanam, and Eso (2002), who approached the problem of incentives differently than other papers in economics. Rather than treating incentives as a constraint that is co-equal with

²Bernheim and Whinston introduced the concept of *coalition-proof* equilibrium as a refinement of Nash equilibrium and focused their analysis on coalition-proof equilibrium outcomes.

³An early example was given by Kelso and Crawford (1982).

other constraints, they searched for the mechanisms to minimize various measures of the distance between the Vickrey–Clarke–Groves payments and the actual payments, given other constraints like efficiency and individual rationality. By varying the distance metric, they varied the maximum gains that bidders can secure by strategically misrepresenting their preferences. This idea, while sure to be controversial because it discards the idea that incentives are constraints, appears promising in practical terms. In my own work, I use a similar idea to create new auction designs, imposing the constraint that outcomes lie in the core with respect to reported preferences.⁴

Before beginning the main analysis, it is worthwhile to ask why the form of the mechanism matters at all. It is a question that has sometimes been asked quite seriously. In 1993, in preparing for the U.S. spectrum auctions, some economists and regulators cited the Coase theorem to argue that the form of the mechanism used could not affect the final allocation, which anyway would be determined efficiently by resale of licenses after the auction. In my public comments, I countered those arguments then by pointing to the U.S. experience with mobile telephones over the preceding decade. In that period, the highly fragmented U.S. market failed to develop the standards and roaming agreements that would make it possible for customers to take their phones from city to city, suggesting that the problem of rationalization was not automatically solved by Coasian bargaining in this market. I also cited the Myerson–Satterthwaite and Vickrey theorems as establishing that the initial allocation can affect the efficiency of the final allocation; Coasian bargaining is not a panacea even in theory.

In theory, introducing a plausible resale option could serve further to limit the set of outcomes that can be achieved by a mechanism designer. Should those constraints be imposed on the analysis? It is easy to find both theoretical and practical reasons why resale might be difficult. The U.S. experience with spectrum trading confirms that building a valuable package by pairwise trades is difficult and time-consuming, leading to large losses of welfare along the way. Yet, eventually, spectrum resources do periodically change hands and these trades have, in reality, allowed the creation of nearly nationwide wireless service networks. In this paper, to focus my analysis on the many issues surrounding the initial trade, I proceed as if resale were impossible.

2. DECENTRALIZATION, MARKET CLEARING, AND THE CORE

The simple example of Table I illustrates the central issues that I will address in this paper. The transaction to be studied involves two goods, labeled A and B, and two bidders or potential buyers, labeled 1 and 2. Suppose initially that there is a single seller. The bidders have quasilinear payoffs and the values for the items shown in Table I.

⁴See Milgrom (2006).

TABLE I

	A	B	AB
1	0	0	12
2	10	10	10

For bidder 1, goods A and B have no value individually, but the bundle or *package* AB has value 12, so the goods are complements. For bidder 2, either good alone is worth 10 and the package is also worth 10, so the goods are substitutes. A pattern like this might arise, for example, in a radio spectrum auction if bidder 1 were a new entrant and bidder 2 were an incumbent mobile telephone carrier. Bidder 1 might need to acquire both spectrum licenses A and B to make it worthwhile to incur the fixed costs of erecting cell sites, establishing an office, and so on. Bidder 2, being an incumbent, may already have incurred those costs and be seeking to add the bandwidth of A or B to its existing capacity to serve a growing set of customers.

The unique efficient allocation in Table I is for bidder 1 to acquire both goods. So, by the first welfare theorem, any market-clearing price vector (p_A, p_B) would have to support that. Then, bidder 1 must prefer to buy both goods and bidder 2 must prefer to buy neither, which requires $p_A + p_B \leq 12$, $p_A \geq 10$ and $p_B \geq 10$. These inequalities cannot be simultaneously satisfied, so *there are no market-clearing prices*.⁵ A glance at the example shows that this is no knife-edge result: the values can be modified considerably without disturbing the nonexistence conclusion.

The biggest surprise in the theory is that the very same conditions that are critical for the existence of competitive equilibrium in Table I are also critical for the satisfactory performance of Vickrey auctions, for the uniqueness of equilibrium in pay-as-bid (“menu”) auctions, and for the ex post equilibrium property of the new ascending proxy auction design that Lawrence Ausubel and I have proposed. To highlight this four-way relationship, let us consider each of the theories in turn.

2.1. *Substitutes and the Existence of Market-Clearing Prices*

Kelso and Crawford (1982) were the first to study a model similar to the package-auction problem, but set in the slightly different context of multiple firms bidding for the services of workers. They showed that when the firms regard workers as substitutes, a certain ascending auction process ends in market-clearing prices. They also displayed an example in which firms do not regard workers as substitutes and competitive equilibrium does not exist.

⁵The nonexistence of competitive equilibrium in similar problems has long been known. Other examples can be found in Kelso and Crawford (1982) and Bikhchandani and Mamer (1997).

Analyses by Gul and Stacchetti (1999) and Milgrom (2000) built on this earlier paper to show for the package-auction model that when goods are substitutes for all the bidders, a competitive equilibrium always exists and that certain auction procedures converge to exact or approximate equilibria.⁶ Strikingly, in a *simultaneous ascending auction* of the kind used by the Federal Communications Commission (FCC) to sell radio spectrum licenses, if bidders behave nonstrategically as price takers during the auction, then the prices converge monotonically up to *approximate* competitive equilibrium prices. The FCC's auction is similar in important respects to the process studied by Kelso and Crawford (1982). It is a collection of several ordinary ascending auctions for spectrum licenses in which bidding for all licenses remains open until there are no new bids for any license.⁷ To see its algorithmic properties, suppose that the bidders behave nonstrategically, always bidding for what they would wish to buy at the prevailing prices. At any time during the auction, the standing high bidder on some particular spectrum license, say license A, is someone who had bid for license A when the prices of other licenses were weakly lower. If the licenses are substitutes, it follows that the bidder still demands license A at the current prices. When the simultaneous ascending auction ends, there is one standing high bidder for each item and no bidder who is willing to bid one increment more for any license. In that sense, the final prices in the auction are approximately market-clearing prices.

There are partial converses to the result that when goods are substitutes, a competitive equilibrium exists. The converses are formulated in terms of the kinds of *environments* in which equilibria can be *guaranteed* to exist. The intended interpretation is that for environments that are reasonably broad and for which it is possible that some bidders have values that are not substitutes, there are value profiles for which market-clearing prices do not exist.

The formal structure that recurs in the theories treated here is the following. Let G denote the set of goods for sale and let $V_{\text{all}} = \{v : 2^G \rightarrow \mathbb{R}_+, v(\emptyset) = 0\}$

⁶Gul and Stacchetti introduced an algorithm that, under certain assumptions, computes to an exact competitive equilibrium.

⁷The FCC's simultaneous ascending auction mechanism consists essentially of a collection of N ascending auctions, one for each item being sold. All take place simultaneously in a series of round. Bidders can add new bids for any item until there is a round or two in which there are no new bids for any item, that is, bidding for all the items closes simultaneously. At the end of each round, there is a provisionally winning bid for each item (the initial provisionally winning bid belongs to the auctioneer) and new bids for a license must exceed the provisional winner by a specified increment. Another important rule in the auction is the Milgrom–Wilson *activity rule*, which utilizes a quantity measure for spectrum licenses (based on the population of the covered geography times the license bandwidth). A bidder's *activity* at a round is the quantity of licenses on which it is the provisionally winning bidder plus the quantity of other licenses on which it makes new bids. A bidder is not permitted to be active on more quantity than its *eligibility*. Bidder i 's eligibility for the first round is set by its cash deposit before the auction. For later rounds, it is set by $E_{n+1}^i = \min(E_n^i, \alpha_n A_{n+1}^i)$, where E_n^i is i 's eligibility at round n , A_n^i is i 's activity in round n , and $\alpha_n \leq 1$ is a parameter set by the auctioneer.

denote the set of all possible values for packages. An *environment* is a subset $V \subset V_{\text{all}}$ of values that bidders in the model might possibly have.

Three environments play important roles in the theory. First is the set of *substitutes* values, V_{sub} , which consists of the values for which the goods are substitutes in the standard sense that raising the price of any one good does not decrease demand for any other good.⁸ Second is the set of *additive values*, V_{add} , which consists of the values v with the property that for any package $x \subset G$, $v(x) = \sum_{g \in x} v(\{g\})$. These are the values for which the value of any package is the sum of the stand-alone value of the constituent items. Last are the *singleton values*, V_1 , which consist of the values v with the property that for every nonempty package $x \subset G$, $v(x) = \max_{g \in x} v(\{g\})$, so that the buyer has no use for more than one good.

Using this notation, the existence theorem just described can be stated as follows: if $V \subset V_{\text{sub}}$, then for any number of bidders n and any value profile drawn from V^n , a competitive equilibrium exists. There are two partial converses. Mine (Milgrom (2000, 2004)) states that if $V_{\text{add}} \subset V$ and $V \not\subset V_{\text{sub}}$, then for any $n \geq 3$, there exists a value profile in V^n such that no competitive equilibrium exists. Another (Gul and Stacchetti (1999)) states that if $V_1 \subset V$ and $V \not\subset V_{\text{sub}}$, then there exists some n and a value profile in V^n such that no competitive equilibrium exists.⁹

These converses reflect serious failings in the performance of the simultaneous ascending auction when goods are not substitutes for some bidder. Suppose, for example, that a bidder who does not regard items A and B as substitutes bids for both of them early in the simultaneous ascending auction when all prices are relatively low. The bidder runs the risk that there will be no more bids for A but that bidding for B will be active, driving its price up to a high level. If A and B are complements for the bidder in the relevant range of prices, then the bidder will no longer want to buy item A, but may be stuck with it under the rules of the auction. This risk—that a bidder may be required to buy an item it no longer wants after other prices have adjusted—is called the *exposure problem*. It poses a dilemma for bidders, who must decide whether to bid aggressively despite the risk involved or whether to eschew bidding for A and B even though the bidder would want to buy the items at the current auction prices. No matter what strategy the bidder adopts, there is a risk that the

⁸In this discrete case, demand is expressed by a correspondence $D(p)$, not a function. Kelso and Crawford (1982) extend the definition of substitutes to this case as follows. Demand satisfies the *substitutes* condition if for any price vectors $p \geq p'$ and any package $x \in D(p)$, there is some $x' \in D(p')$ such that for all j with $p_j = p'_j$, we have $x_j \geq x'_j$. When goods are divisible and D is singleton-valued, this corresponds exactly to the familiar definition.

⁹The requirements that $V_{\text{add}} \subset V$ or $V_1 \subset V$ is not dispensable. Sun and Yang (2006) identified an alternative formulation for which equilibrium exists. In their formulation, there are two types of goods. Goods of each type are substitutes for one another but are complements for goods of the other type, as in the example of left shoes and right shoes.

resulting allocation will be quite inefficient, in sharp contrast to the situation when goods are substitutes.

One can try to rescue competitive equilibrium theory for this situation by extending the space of prices beyond the anonymous, linear prices of neoclassical price theory. The revised theory would employ prices that are either nonlinear, so that the price of a package of goods may be different from the sum of the item prices, or nonanonymous, so that different bidders face different prices for the same packages, or both. It is not hard to show that with nonlinear and nonanonymous prices, one can indeed support efficient allocations. The relevant definitions of competitive equilibrium with such prices and the properties of equilibrium have been explored by various authors (Bikhchandani and Ostroy (2002), de Vries and Vohra (2003), Parkes and Ungar (2002)).

The competitive equilibrium theory that we have been discussing abstracts from strategic behavior, which is an important omission for the applications we wish to study. Another relevant tradition in economic theory is that of *mechanism design*, which treats incentives as constraints that are binding in the same sense as technical and resource constraints. I will argue subsequently that the very same substitutes condition that is critical for the existence of competitive equilibrium is also critical for satisfactory performance of certain important incentive mechanisms.

The two best known theoretical mechanisms that employ nonanonymous package prices are the *pay-as-bid package auction* or *menu auction* (Bernheim and Whinston (1986)) and the *Vickrey auction* or *pivot mechanism* (Vickrey (1961), Clarke (1971), Groves (1973)). Both auctions select the outcome that maximizes the total winning bid. The pay-as-bid package auction is the mechanism in which each winning bidder pays the amount of its own winning bid for the items it acquires. The Vickrey auction uses a more complex pricing formula, in which the payment by each bidder j is computed so that the total payoff of the seller and the bidders other than j is the same as if bidder j had not participated in the auction.

The pay-as-bid design is used in practice in an increasing number of applications, but I am unaware of even a single application of the multi-item Vickrey auction with two or more *different kinds* of goods for sale. Our next task is to understand the reason for this observed pattern by evaluating the strengths and weaknesses of these two mechanisms as practical, real-world designs.

Substitutes, the Core, and the Vickrey Mechanism

Is the core a relevant standard?

The *core* is the set of *allocations* of goods and money with the property that no coalition can do better by trading on its own.¹⁰ One can alternatively identify

¹⁰Two slightly different versions of the core can be defined, depending on the precise meaning given to the phrase “do better,” which can refer to either weak or strict Pareto improvement for the coalition.

the core with the set of *payoff profiles* that correspond to the core allocations. These payoff profiles are often called *imputations*, because the payoffs are imputed from the underlying allocation.

The core is an important concept for applied mechanism design for several reasons. First, if a proposed mechanism generates outcomes outside the core and if the players are well informed, then some *coalition* of the participants can do better by rejecting the proposed mechanism and allocation in favor of something else, undermining the viability of the mechanism. In the particular case of an auction, if the outcome is not in the core, then there is a set of bidders that could profitably offer a better deal to the seller who would benefit by accepting it. This observation introduces the second idea, namely, that an allocation is in the core if and only if all the players' payoffs are "competitive" in a sense that is useful especially for public sector applications, where competitive outcomes are sought.

To define the core in this context, we must first describe the relevant coalitional game. Let N denote the set of players in the game, including both the seller and the bidders. Let X denote the set of feasible allocations of the goods with typical element x , let x_j denote the goods allocated to player j , let $u_j(x_j)$ denote the value of j 's allocation, and assume "free disposal" so that each u_j is weakly increasing. Given any coalition $S \subset N$, we construct the value $w(S)$ of the coalition as the maximum value of the allocation of the goods among the members of S :

$$(1) \quad w(S) \equiv \begin{cases} 0, & \text{if seller} \notin S, \\ \max_{x \in X} \sum_{j \in S} u_j(x_j), & \text{if seller} \in S. \end{cases}$$

If the seller is not in the coalition, there are no goods to distribute. If the seller is in the coalition, then the goods are distributed to maximize total value. Implicitly, we assume that the total payoff can be arbitrarily reallocated by the members of coalition S among themselves using side payments.

With this definition of the coalitional value function w , the core of this transferable utility game is

$$(2) \quad \text{core}(N, w) = \left\{ \pi \geq 0 \mid \sum_{j \in N} \pi_j = w(N), (\forall S) \sum_{j \in S} \pi_j \geq w(S) \right\}.$$

Notice that the core of a game with a single seller is always nonempty, because it includes the imputation at which the seller gets the whole value $w(N)$ and each bidder gets zero.

The idea that the core formalizes a notion of price competition is as old as mathematical economics itself: Walras made the first showing that any competitive equilibrium allocation is a core allocation. For the package-auction model with transferable utility, there is a second tight connection to reinforce the idea

that core payoffs are *competitive*: an imputation π is in the core if and only if it is a vector of competitive equilibrium prices in a certain production economy derived from (N, w) . The agents in this economy include the seller and the bidders who compose N . Any agent in this economy can hire other agents to form a firm with members $S \subseteq N$ whose gross revenue is the coalition's value $w(S)$. Suppose π is a vector of wages for the agents in this economy. Then the net profit to an agent who organizes a firm with members S is $w(S) - \sum_{j \in S} \pi_j$. If a competitive equilibrium exists for this economy, then the outcome must be efficient and that can be achieved by creating a firm from N , the coalition of the whole. Equilibrium requires zero profits for the agent who organizes the firm, so $w(N) - \sum_{j \in N} \pi_j = 0$. Also, there must be no way for any agent to make a positive profit by forming a different coalition S , so $w(S) - \sum_{j \in S} \pi_j \leq 0$. These are exactly the conditions that $\pi \in \text{core}(N, w)$, so π is a competitive equilibrium price vector for the economy if and only if it is a core imputation.

With a single seller, the core is always nonempty: the efficient allocation that gives all the surplus to the seller and zero to each bidder is a core allocation. However, the core can be empty when there are multiple sellers. For example, if we reinterpret Table I to describe a situation in which items A and B belong to different sellers, then the core is empty. To see this, denote the sellers of A and B by the corresponding lowercase letters, so the set of players is $N = \{a, b, 1, 2\}$ and suppose that π is a core imputation. We obtain a contradiction as follows. Because $w(N) = w(N - \{2\}) = 12$, we must have $\pi_2 = 0$. Hence, by the core inequalities, $\pi_a = \pi_a + \pi_2 \geq w(a2) = 10$ and symmetrically $\pi_b \geq 10$. Finally, $\pi_1 \geq 0$. Adding these up, we get $w(N) = \pi_a + \pi_b + \pi_1 + \pi_2 \geq 20$, but $w(N) = 12$, which is the promised contradiction.

The emptiness of the core with multiple sellers makes it impossible to create any mechanism that yields "competitive" payoffs even when other incentive problems are absent. It implies that no matter what mechanism a designer might propose, if the participants are well enough informed, there is always some coalition among them that can do better by refusing to cooperate in favor of negotiating on their own. Thus, voluntary multilateral exchanges with universal participation may be impossible to organize, even apart from the efficiency and incentive problems. They pose a very tough challenge for practical mechanism designers.

Even in cases where the core is nonempty, core allocations cannot always be implemented in an incentive-compatible way. Next, we study a particularly famous incentive-compatible mechanism to analyze when to select core allocations and what happens.

Vickrey package auctions

The multi-item Vickrey auction has attracted more academic attention than any other mechanism because of its remarkable theoretical properties. It is well known that in the Vickrey auction design, (i) it is a dominant strategy

TABLE II

	A	B	AB
1	0	0	12
2	10	10	10
3	10	10	10

for bidders to report their value information truthfully and, when they do so, (ii) the mechanism implements an allocation that maximizes total value. Also, (iii) losing bidders in the Vickrey auction design pay zero. Green and Laffont (1977) showed that if the set of bidders' values includes all possible values, then the Vickrey auction is the unique mechanism with the preceding properties (i)–(iii). Holmstrom (1979) improved that result in an important way, showing that the Green–Laffont conclusion is true even if the set of bidder values is restricted, provided the set is path-connected.

In the example of Table I, a Vickrey auction would assign both goods to bidder 1 at a price of 10. The bidders in the example have dominant strategies and the outcome is efficient, but there is also an additional property: the outcome is a core allocation for the three player game that involves the two bidders and the seller. Given this favorable performance, why is the Vickrey package auction never used?

A critical problem with the Vickrey mechanism is that even when there is just a single seller (which guarantees that the core is nonempty), the mechanism may not select core allocations. To illustrate the problem, let us introduce a third bidder with the same values as bidder 2, as shown in Table II.

In this case, the Vickrey outcome assigns the two items to bidders 2 and 3. At the efficient solution without bidder 3, bidders 1 and 2 would acquire goods that they value at 12. With bidder 3, only one good is allocated to them and they get a value of 10. The difference is $12 - 10 = 2$, and so that is the price that bidder 3 pays. By symmetry, bidder 2 pays the same. The outcome is efficient and the seller's revenue is 4.

The problem is that these payoffs are not competitive. At any core imputation in this example, the losing bidder 1 must get zero and the seller, whom we denote as player 0, must receive at least 12. If the seller's revenue is less than 12, then the coalition $\{0, 1\}$ that consists of the seller and bidder 1 could profitably defect and divide $w(01) = 12$ among themselves. The Vickrey outcome lies outside the core in this example because the seller's revenue is too low. Intuitively, the revenue shortfall appears severe, but the example is not extreme. In the variation of this example in which the number 12 in Table II is replaced by 10, the prices are zero and seller's Vickrey revenue is zero, despite the fact that all three bidders are willing to pay 10 for the package.

This example is typical of the general way in which a Vickrey auction can lead to uncompetitive payoffs: *the seller's revenue is too low*. Whenever the Vickrey

payoff is not in the core, the seller's revenue is always uncompetitively low—strictly less than the smallest seller revenue at any core allocation—and each bidder's payoff is as high as possible consistent with the competitive standards.

These points are easy to prove, so let us prove them. Recall that the Vickrey outcome is always efficient, so the total payoff of all the participants is $w(N)$. Recall, too, that bidder j 's Vickrey payments are computed so that the total payoff of the other players, including the seller, is $w(N - j)$, just as if j did not participate, regardless of j 's report. Hence, j 's Vickrey payoff must be $\bar{\pi}_j = w(N) - w(N - j)$. If for some proposed payoff allocation π it were true that $\pi_j > \bar{\pi}_j$, then the total payoff to coalition $N - j$ could be at most $w(N) - \pi_j < w(N - j)$, so $\pi \notin \text{core}(N, w)$. Hence, $\bar{\pi}_j \geq \max\{\pi_j | \pi \in \text{core}(N, w)\}$. Moreover, the imputation with payoffs $\bar{\pi}_j$ for player j , $w(N - j)$ for the seller, and zero for the remaining bidders is a core imputation. So, we have proved the theorem of Ausubel and Milgrom (2002) and Bikhchandani and Ostroy (2002) that $\bar{\pi}_j = \max\{\pi_j | \pi \in \text{core}(N, w)\}$: each individual bidder's Vickrey payoff is its highest payoff among core imputations. Also, if π is a core imputation, then $\pi_0 = w(N) - \sum_{j \in N-0} \pi_j \geq w(N) - \sum_{j \in N-0} \bar{\pi}_j = \bar{\pi}_0$. Hence, $\bar{\pi}_0 \leq \min\{\pi_0 | \pi \in \text{core}(N, w)\}$, with strict inequality whenever $\bar{\pi} \notin \text{core}(N, w)$.

So, we have established that the Vickrey outcome lies outside the core exactly when the seller's Vickrey revenues are less than its minimum revenues for core payoff and that such low Vickrey revenue outcomes are possible. Yet, given the considerable advantages of the Vickrey mechanism, that is not quite the end of the story. Is the problem of low revenues really so pervasive as to overcome the advantages of the Vickrey design? Are the low-revenue outcomes in some sense common? On the contrary, is there a large and important set of environments in which low-revenue outcomes are known to be excluded? We pose these questions in an analytically tractable form by asking the following question: In what sets of environments is the Vickrey payoff vector $\bar{\pi}$ in the core?

This question can be studied at two different analytical levels. One level operates in the traditional fashion of cooperative game theory, treating the *coalition-value* function w as the primitive that describes the relevant environments. The second level is the one commonly adopted for analyses of resource allocation problems, in which the possible environments are described by the set V of *goods-valuation* functions that the bidders might have. For both formulations, we adopt the same general approach that we applied to our study of the existence of competitive equilibrium: given a restriction on a set of possible profiles, our objective is to determine whether the Vickrey outcome is in the core for all profiles in the relevant environment. The following results of this form are drawn from Ausubel and Milgrom (2002).

The traditional cooperative game formulation treats the coalition value function w (usually called the characteristic function) as primitive. We will say that the function w is *bidder-submodular* if for all coalitions S that include the seller and player j , $w(S) - w(S - j)$ does not increase as bidders are added

to S . Now, suppose that a set of potential bidders is exogenously identified but that only some of the bidders will actually wind up participating in the auction. The Vickrey auction payoffs will then depend on the set of players S that participate. Denote the corresponding Vickrey payoff vector by $\bar{\pi}(S)$ and denote the core of the restricted game by $\text{core}(S, w)$.

When can we be sure that the Vickrey payoff vector will be in the core regardless of which bidders participate? The formal answer is that $\bar{\pi}(S) \in \text{core}(S, w)$ for every $S \subset N$ if and only if π is bidder-submodular. Notice from our preceding analysis that $\bar{\pi}_j(S) = w(S) - w(S - j)$. If this increases when S grows to $S + j'$, then $\bar{\pi}(S + j')$ is blocked by $S - j$, so bidder submodularity is necessary. The bidder-submodularity condition implies that when the auction participants are the members of coalition S , the Vickrey payoff for any $j \in S$ is smaller than j 's marginal contribution to any subcoalition $T \subset S$. Consequently, any subcoalition T will prefer to include j and pay j 's Vickrey payoff rather than to exclude j . An argument along those lines proves that when the condition holds, no coalition blocks, so $\bar{\pi}(S)$ is a core imputation and hence entails a competitive payoff to the seller.¹¹

In standard economic analysis, individual preferences are the primitives from which all other values are derived. For example, the worth of any coalition is derived from individual preferences by the maximization problem (1). Classes of environments can be described by restrictions on the set of individual preferences, such as the common restriction in competitive equilibrium theory that preferences are convex. In our problem, individual preferences are described by the values that bidders may put on various packages, so our second formulation takes the set of possible value functions V as the relevant primitive.

Given n bidders, a profile of values $v \in V^n$, and a corresponding coalitional value function w , the Vickrey payoff for any bidder j is given by $\bar{\pi}_j(v) = w(N) - w(N - j)$ and for the seller by $\bar{\pi}_0(v) = w(N) - \sum_{j \in N-0} \bar{\pi}_j(v)$. The theorems that emerge from this analysis are identical in form to the previously described theorems about the existence of competitive equilibrium. First, if $V \subset V_{\text{sub}}$, then for every n and $v \in V^n$, $\bar{\pi}(v) \in \text{core}(N, w)$. In words, if the goods are substitutes for all the bidders, then the Vickrey payoffs are in the core. Second, if $V_{\text{add}} \subset V$ and $V \not\subset V_{\text{sub}}$, then for every $n \geq 4$, there is a profile $v \in V^n$ such that $\bar{\pi}(v) \notin \text{core}(N, w)$.

If we compare these theorems about Vickrey outcomes lying in the core with the earlier results about existence of competitive equilibrium, a consistent theme emerges. Environments in which all goods are substitutes ($V \subset V_{\text{sub}}$) are ones in which decentralized mechanisms have the identified desirable properties. For any superset of the substitutes environments ($V_{\text{sub}} \subset V \neq V_{\text{sub}}$), the

¹¹Bikhchandani and Ostroy (2002) independently proved that if π is bidder submodular, then the Vickrey outcome $\bar{\pi}(N)$ is in the core.

properties fail: the set V_{sub} is a maximal set on which “satisfactory performance” is guaranteed. Moreover, if we constrain the set of environments V to be sufficiently inclusive, so that V includes either the additive valuations ($V_{\text{add}} \subset V$) or the singleton valuations ($V_1 \subset V$), then the desired properties are guaranteed *if and only if* V is limited to substitutes valuations ($V \subset V_{\text{sub}}$).

In pay-as-bid (“menu”) auctions, prices go up as participation increases or as bids increase, but the same is not true in the Vickrey design: revenues decline as we move from the situation in Table I to that described in Table II. This is problematic for several reasons. First, if the situation in Table II prevailed, the seller would have an incentive to restrict participation by bidder 3 or to seek to disqualify its bid after the auction, because that would restore Table I and raise the auction price. This can sometimes be a serious issue. It is not uncommon in large asset sales for bidders’ qualifications to be reviewed most carefully after the bids are opened and tentative winners are identified. Bids in large asset sales are often submitted with explicit or implicit contingencies, becoming fully effective only if approved by government regulators, unions, creditors, and so on, and the seller may reasonably reject bids if it judges that there is too much risk that the deal will not actually be closed. Clearly, such a procedure would be problematic for the Vickrey design.

A second, related problem is that bidder 2 might seek to transform the situation in Table I into the one in Table II by bidding under two different names, as bidders 2 and 3. By that device, bidder 2 would win both items at a total cost of 4, which is less than its value of 10.¹²

Third, suppose that the situation is similar to the one described in Table II, but with bidder 1’s value raised from 12 to 24. In this situation, it is efficient for bidder 1 to win and that is the outcome if the bidders play their dominant strategies. Bidders 2 and 3, however, have a profitable joint deviation in which they report that their values are 20 instead of 10. By playing those strategies, each wins an item and pays a price of 4, so this joint deviation is profitable for the deviators. This is an example in which losing bidders can profitably collude to become winning bidders.

None of these disadvantages is shared by the more popular pay-as-bid auction mechanism. In that design, higher bids result in increased revenues, a bidder cannot reduce the price it must pay to win a package by splitting its bids among two or more names, and losing bidders cannot collude profitably without including some winning bidders in their coalition.

Are these revenue and strategic problems really the reasons that multi-item Vickrey auctions are never seen in practice? There are two kinds of objections that one might raise to this explanation. The first holds that, at least in government-run auctions, revenues should not be the primary objective. The

¹²Yokoo, Sakurai, and Matsubara (2004) emphasized that this problem is particularly troubling for Vickrey auctions run at online auction sites, because bidders’ identities online are notoriously hard to verify.

second is that there are other characteristics of the Vickrey auction design that account for its unpopularity.

The first objection holds that in government run auctions, *efficiency* and not revenues should be the primary objective. In U.S. spectrum auctions, efficiency is indeed one goal of the auction program, but it is not the only one. The Communications Act 309(j)(3) mandates both efficiency and revenue goals, as well as other goals.¹³ In practice, in U.S. spectrum auctions, revenue has come to be a more central concern than efficiency, perhaps because efficiency is so hard to measure in real auctions. Moreover, even for public sector auctions, there are strong arguments to be made that revenues *ought* to be a critical concern. Revenues earned at auction can, in principle, be offset by lower taxes, reducing the distortionary effects of taxation on the economy.¹⁴ Moreover, the auction is only part of the resource allocation game. If public assets are to be sold for less than their market value, that creates rents and encourages wasteful expenditures on rent-seeking behavior. For both of these reasons, even beneficent public decision makers whose main concern is efficiency would still pay close attention to the level of revenues that an auction generates.

The other possible objections to our analysis focus on alternative explanations of the absence of Vickrey auctions in practice. In their aptly titled paper, “Why Are Vickrey Auctions Rare?,” Rothkopf, Teisberg, and Kahn (1990) proposed one such alternative. They argued that bidders will be reluctant to report their actual values to the auctioneer, for fear that the report may eventually be used against them in other transactions. For example, a bidder for a supply contract may want to keep his/her cost information secret in case there are change orders that require terms to be renegotiated or in case there is follow-up work for this customer or others. According to this theory, it is not only multi-item applications of the Vickrey auction that should be rare, but all Vickrey auctions and related designs, in which bidders reveal the maximum price they are willing to pay. The evidence does not support that contention. Some real bidders do report maximum bids to proxy bidders millions of times per day at electronic auction markets like eBay, although others do not (Ockenfels and Roth (2003)). All of the bids to sell the advertisements that Google places beside its search listings are based on values reported privately to the auctioneer. The Rothkopf–Teisberg–Kahn theory points to a consideration that is important for some applications, but it cannot explain the frequent use of Vickrey-like protocols for online single-item auctions but never for multi-item auctions.

A second frequently heard alternative explanation is that, for multi-item auctions, the Vickrey design is too computationally and cognitively demanding,

¹³These include the promoting development and rapid deployment of new technologies, products and services, promoting competition, ensuring the wide availability of new technologies, avoiding excessive concentration of licenses, and disseminating licenses among a variety of applicants (specifically including small businesses and rural telephone companies).

¹⁴Klemperer (2000) emphasized this point.

because there are so many possible packages for which bidders must report values. This, too, is a genuine concern for some auctions, particularly those with many distinct items for sale. Complexity was one of the principal considerations in designing the details of the FCC's spectrum auctions and it continues to be a principal consideration in proposals for new designs. Nevertheless, complexity cannot explain the absence of Vickrey auctions even when there are a small number of items for sale and, in particular, when there are just two items, because the cognitive and computational demands in such cases are low. Moreover, any theory of the absence of Vickrey auctions based on its cognitive and computational burdens ultimately stumbles on the fact that other sealed-bid package-auction designs are used in industrial applications. Like the Vickrey auction, these sealed-bid package designs require that bidders name prices for each package of interest to them. Also, like the Vickrey auction, they require that the auctioneer solve a combinatorial optimization problem to identify the winning bids. It is far from clear why these designs are possible, but the Vickrey design is not.

So, why are multi-item Vickrey auctions rare? The revenue reasons and related monotonicity problems are sufficiently pervasive and serious to yield a plausible answer, while alternative theories, consistently applied, explain both too much and not enough. Identifying the real reasons for the absence of Vickrey designs is important. If new auction designs are to be adopted, they will need to avoid the problems of the Vickrey design and to match the benefits of the best existing designs.

2.2. Menu Auctions

The package auctions that have been used most frequently in practice are known variously among scholars as *menu auctions* or *pay-as-bid package auctions*. Although the detailed rules for practical versions of these auctions vary greatly among applications, it is useful for theorizing to study a canonical form in which a bidder j may submit multiple bids of the form (x_{jk}, b_{jk}) , $k = 1, \dots, K$, where x_{jk} denotes a package of items and b_{jk} denotes the amount bid for the package. In our model, the seller may accept only one bid from each bidder and may impose a variety of constraints on the outcomes. For example, in government sponsored auctions, it may require that some fraction of the sales go to minority-owned bidders or that some measure of the concentration of sales be below some threshold. Commercial buyers may add other constraints to the allocation. For example, they may seek to avoid supply interruptions by ensuring that suppliers have sufficient additional capacity among them to meet potential increases in demand or by requiring geographic diversity of supply, so that strikes, hurricanes, earthquakes, and such would not completely cut off essential supplies. There may be lower bounds on the fraction of business won by bidders with high quality ratings. Long-term contracts may guarantee some suppliers a certain minimum fraction of the business on

the contracted terms. Subject to these and other constraints, the seller accepts the combination of bids that maximizes its revenue or, in the case of procurement, minimizes its costs.

This canonical form provides a helpful approximation of several actual auctions. In the public sector, the city of London uses a kind of first-price package auction to acquire contracts for groups of bus routes (Cantillon and Pesendorfer (2003)). In Chile, milk supplies to schools are acquired in a similar way (Epstein, Henríquez, Catalán, Weintraub, and Martínez (2002)). Business firms, including IBM (Hohner, Rich, Ng, Reed, Davenport, Kalagnanam, Lee, and An (2003)), routinely use such auctions and several companies offer software to support these auctions.

Several of these applications could naturally involve values like those in Table I. For example, consider the problem of contracting for London bus routes. Suppose a group of contracts expires in one part of the city and new contracts are put out to bid. Bidder 1 might be a bus company that has no garage in that part of the city and could serve the routes only by building a garage there. The fixed costs and associated scale economies do not make it worthwhile to introduce service there unless it can serve several bus routes. Bidder 2 could be a bus company that has an existing garage in that part of the city, but whose garage has limited additional capacity. It would then find it economical to bid on some but not all of the routes. Value and cost patterns of this kind are common in any industry where there are significant capacity limits and fixed costs to build new capacity.

The pay-as-bid package/menu auction has been analyzed by Bernheim and Whinston (1986) and, again, related to core allocations. Let us say that a core allocation is *bidder optimal* if there is no other core allocation that all bidders weakly prefer and some bidder strictly prefers. The full-information, coalition-proof equilibrium allocations of the auction coincide with the bidder-optimal core allocations.

An intuition for this result can be gotten by referring to the special case of this model in which only one good is sold and the seller imposes no constraints. In this case, there is no question about what package a bid is for, so a bid corresponds to a single number. The model in this case coincides with the familiar Bertrand auction model. If the highest value for the single good is 10 and the second highest value is 8, then the Bertrand equilibrium outcome is that the highest valuing bidder acquires the item for a price of 8. If the price were any lower, then the seller could form a coalition with the bidder having the second highest value for their mutual benefit, so 2 is the highest payoff that the winning bidder can earn in the core. At any core allocation, losers must get zero, because the coalition of the winner and the seller must get the total value of 10. The Bertrand outcome is the core allocation with the highest possible payoff for every bidder.

In general, given a bidder-optimal imputation π , the corresponding strategies are for each player j to place the mutually exclusive bids $(x_{jk}, b_j(x_{jk}))$

for every possible package x_{jk} , where $b_j(x_{jk}) = v_j(x_{jk}) - \pi_j$. According to this strategy, the amount bid for each package is equal to the bidder's actual value of the package marked down by a profit target, π_j , which is the same for every package. Consequently, if any of its bids is winning, then j will earn a profit of exactly π_j . At the coalition-proof equilibrium, that vector of profit targets π_j is a bidder-optimal core imputation.

From a modern perspective, this analysis is not completely satisfying. The full-information model that underlies the analysis of the pay-as-bid auction is contrary to the main theme of mechanism design theory, which emphasizes that bidders have private information about their values. Moreover, even if we accept both the assumption of full information and the coalition-proof solution concept as helpful simplifications, we are still left with a multiple equilibrium problem, because there may be many bidder-optimal core imputations π . It would appear that the best case for this analysis of the pay-as-bid auction is the one where the bidder-optimal core allocation is unique.

We have previously observed that the best core outcome for bidder j pays $\bar{\pi}_j$, j 's Vickrey payoff. It follows that there is a unique bidder-optimal core imputation *if and only if* the Vickrey payoff is in the core— $\bar{\pi} \in \text{core}(N, w)$ —and in that case the pay-as-bid auction equilibrium payoff coincides with the Vickrey payoff.

On what class of environments V can we guarantee that the bidder-optimal core imputation is unique? If the seller imposes no constraints on packages besides the constraint that each good is sold just once, then the answer follows directly from the theorems about Vickrey payoffs. Suppose that $V_{\text{add}} \subset V$ or that $V_s \subset V$. Then there is a unique bidder-optimal core imputation (the Vickrey imputation) for every profile from V^n if and only if $V \subset V_{\text{sub}}$, that is, if and only if all possible values regard the goods as substitutes.

This review of the pay-as-bid auction leaves the appeal of this design somewhat obscure. The known results depend on strong information assumptions that are unlikely to be helpful in reality. Interestingly, even then, with the complete information and equilibrium selection assumptions, the theory makes a specific prediction only for the case where goods are substitutes.

Probably the best way to understand the popularity of the pay-as-bid design is to view it not as the result of a conscious process of auction design, but as the natural evolution of older and simpler industrial procurement practices. As computing technologies improved, some industrial buyers were led to use them to minimize procurement costs while meeting business requirements, using posted and negotiated prices as data for the optimization. Sellers, aware of their customers' practices, would be led to quote prices strategically. When sellers further become aware of each others' circumstances, what has evolved is the pay-as-bid package auction.

What our review makes clear is that whatever the merits of these standard auction forms—the simultaneous ascending auction, the Vickrey auction, and the pay-as-bid package auction—they are far from satisfactory mechanisms for

environments in which goods are not substitutes. There is a need for new theories and new mechanisms, and there is an effort underway to fill that need.

3. EXPERIMENTAL DESIGNS AND RELATED THEORY

Despite the theoretical difficulty of constructing workable mechanisms when goods are not substitutes, some mechanisms seem to work well in economic experiments. Although the experiments to date are hardly conclusive, they are provocative and suggest what may be useful new directions for both theoretical and practical mechanism design. We briefly review two of these experiments and the directions in which they point.

The first is a package-auction experiment that was performed under a contract between Cybernomics and the U.S. Federal Communications Commission to compare two kinds of auctions for selling radio spectrum licenses: the FCC's regular *simultaneous ascending auction* and a related package auction. In the package-auction variation, just as in the Vickrey and pay-as-bid package auctions described earlier, a bidder who wants to acquire a package of licenses, say ABC, can bid a single price for that package and be assured that it will acquire all licenses in the package or none—never just part of the package. Unlike the one-shot package auctions we discussed earlier, bidding in both of the experimental auctions takes place in a series of rounds. After each round, bidders are informed about the standing of their bids and can submit new, higher bids as well as bids for different packages.

Cybernomics presented its results to the FCC in a report and at a conference, where they were represented by two highly regarded academic experimenters: Vernon Smith and David Porter. Details of the experiment, which are quite important for assessing the results, are described in Cybernomics (2000).

The environments tested in the experiments included one with additive values only and several with *synergies*—that is, complementarities—among geographically adjacent spectrum licenses. The experimenters reported mainly about the efficiency of the auction outcomes, as measured by the ratio of the value of the allocation achieved divided by the maximum possible value. Their summary results are very striking. With no synergies, efficiency is high for both the simultaneous ascending auction and the ascending package auction. As geographical synergies grow stronger, the efficiency measure for the simultaneous ascending auction falls from 97% to 79%, while the efficiency measure for the package auction falls from 99% to 96%. The suggested interpretation is that package auctions of this sort overcome the exposure problem and permit the mechanism to identify efficient outcomes.

The Cybernomics report is not detailed enough to enable a fully satisfactory assessment of its results. The FCC contract did not require that detailed experimental data be turned over to the sponsors. When the FCC and I later asked for the data, we were told that they had been lost and cannot be recovered. I have had similar problems obtaining data for several other package

auction experiments. The unavailability of detailed data about bidder values and bids submitted makes it impossible to test important hypotheses about bidder behavior, to evaluate the complexity of the choice problems facing the bidders, or to compare the bidder payoffs and auction revenues with a standard such as the core. The Cybernomics report does find that revenues were lower and sometimes much lower for the package auction treatment than for the corresponding simultaneous ascending auction without package bidding. The missing data also make it impossible to compute measures of efficiency that are more comparable across experimental settings. The measure actually used has the undesirable property that it is increased when all package values are increased by 1 per item. A better measure might reduce the numerator and denominator by the value of the random allocation.

The second experiment that I will describe is older and, on the surface, appears to deal with a very different kind of problem. It is one in which the bidders' values for packages are additive (elements of V_{add}). The complexity of the problem in this experiment arises not from the package valuations but from constraints imposed by the seller about the combinations that can be sold.

The resource to be allocated in this experiment is a notional segment of railroad track in Sweden. In the experiment, a single north–south track segment is to be scheduled for use by fast and slow trains and by northbound and southbound trains. The bidders derive value from *train trips*, each of which is a one-time use of the track in one direction. A trip is characterized by a direction (north or south), a departure time, and an arrival time. It is assumed for the experiment that trains maintain a constant speed during any trip. The auctioneer/scheduler imposes the constraint that the trains maintain a safe distance from one another. Although the values are additive, the resource allocation problem is complicated because the set of feasible schedules, in which no two trains crash or come dangerously close, is derived from a complicated set of constraints.

Can such a seemingly complicated resource allocation problem be guided successfully by an auction? Brewer and Plott (1996) ran an experiment to investigate that possibility. The set of possible trips is illustrated in Figure 1. On the horizontal axis in the figure is time and on the vertical is location. Nine trips are illustrated and labeled A through I, with times that include within them any required margins of safety. A *schedule* is a collection of trips. Schedules are constrained so that no two trips intersect in the diagram, that is, no trains crash.

In the experiment, each of 10 bidders has a randomly drawn value for each of the 9 possible trips. The rules of the auction are simple and easy to understand. During the auction, each bidder can bid a price for any trip from A through I. The state of the auction is described at each moment in time by the highest bid on each of these routes and by the identities of the highest bidders. Whenever the state of the auction changes, the auctioneer computes the feasible schedule that maximizes the total bid. It then publicly posts the vector of current prices




FIGURE 1.—Scheduling diagram for routes in the experimental rail system.

and the new tentative schedule. When a period of time T passes with no new bids, the auction is closed and the current schedule is adopted.

For the trains experiment, the subjects were Caltech undergraduates. Each subject is paid according to the total value of the trips he or she acquires at auction minus the total price paid for those trips.

The reported efficiencies achieved in the experiment were astonishing. In 18 of the 21 trials, the outcome was 100% efficient. For the remaining three trials, the ratio of the value of the outcome to the maximum possible value averaged 97%. It is hard to imagine a better efficiency result with human bidders.

One can certainly nitpick about the experimental details, criticizing how poorly it represents an actual rail system. Real shippers do not care only about single trips, but combinations of trips, or at least about return trips for their equipment. There are also likely to be substitution possibilities among trips. A shipper that prefers to send its cargo train north at 2:00 PM might well be willing to substitute a trip at 2:05 PM or even much later if the cost is sufficiently lower.

For advancing the study of market design, however, the experiment is very helpful in precisely the form that it was conducted. By omitting complexity in the bidders' values, it isolates the effects of complexity that comes from the commonly known scheduling constraints, allowing us to evaluate the effects of those constraints alone. The experiment points to the striking hypothesis that such side constraints are unproblematic for the design of centralized market mechanisms.

Complexity in resource allocation can arise in different ways that may not all have the same implications for mechanism design. This is seen most easily by

setting aside strategic issues and treating the auction as a planning algorithm. In the train scheduling experiment, the *computational complexity* of scheduling might well be overwhelming for the bidders, but is managed adequately by the auctioneer's algorithm. On the other hand, the *communication complexity* of reporting the bidders' preferences is low in this experiment, because each bidder's values for all sets of trips can be represented by just nine numbers—a value for each route—rather than by $2^9 - 1$ numbers—a value for each non-empty package of routes. Given this observation, we might conjecture that the Brewer–Plott mechanism could quickly identify an efficient allocation, while a fuller version in which shippers care about their whole schedules and so value packages might take many more rounds to locate an optimum.

Given that perspective, it is useful to try to gain insight into the findings of the train schedules experiment by first treating the auction as a planning algorithm and asking what happens if the bidders act as price takers. Just as for the simultaneous ascending auction, any bidder at a round who is not a tentative winner for trip k might be taken as regarding the relevant trip price as being one increment higher than the current posted trip price, p_k . Think of the auctioneer as a firm that can “produce” various feasible collections of trips; it produces to maximize the total value of output at the final auction prices. At the end of the auction, each bidder, acting as a price taker, demands the collection that maximizes its own payoff at its relevant prices, which are all within one increment of the final auction prices. Each trip that is produced by the firm is consumed by some bidder, so markets clear. The outcome is an approximate competitive equilibrium in the same sense that was relevant for our study of the simultaneous ascending auction. By the first welfare theorem, the outcome is approximately efficient; any loss of efficiency is bounded in proportion to the bid increment. If the bid increments are small, the outcome is exactly efficient.

Ascending Proxy Auctions

Experiments like these are inspiring for what they achieve, but they leave many of the most important questions unanswered. How could the package auction produce such high efficiency for the FCC experiment? Are there generalizable lessons there? Are there mechanisms that can simultaneously deal with complex valuations such as those used in the FCC–Cybernomics experiment and complex constraints such as those of the Brewer–Plott railroad scheduling experiment? Do the results depend on the ignorance and unsophistication of bidders, as suggested by the analysis based on price-taking behavior? Can we go beyond the assumption of price-taking behavior to incorporate strategic incentives into the analysis?

To evaluate these questions, we will consider an auction model introduced by Ausubel and Milgrom (2002) that allows package bids and constraints of many kinds. This model, which incorporates a bidding algorithm, has a dual

interpretation. When we analyze it as a model of the laboratory findings, we interpret it as an extensive form model of how the preceding auctions might be played by naïve bidders. When we analyze it strategically, we interpret it as a direct revelation mechanism in which bidders report values or preferences and the algorithm is part of the definition of the mechanism. In the latter case, reporting incentives are analyzed in terms of game-theoretic solutions.

In the model, each of the bidders $j = 1, \dots, n$ has a finite set of feasible bids X_j and a strict ordering \succ_j over this set. Bids in X_j are interpreted to be mutually exclusive offers. For the applications described in the preceding text, we may interpret X_j narrowly as a set of pairs (x_j, p_j) that describe a package of goods (or a train trip) and a price, and interpret \succ_j as the preference ranking determined by the bidder's net profit. The abstract formal treatment, however, allows other important interpretations. For example, it encompasses the possibility that bidders are budget constrained or, in a procurement auction, that bidders can vary the quality characteristics of their goods.

We assume that each set X_j includes a *null bid* \emptyset_j that yields the bidder a payoff of zero, and is interpreted to mean that the bidder is not a "winner" and is excluded from the auction outcome. For simplicity, let us edit X_j to exclude any bids that are less preferred than \emptyset_j .

For the seller, the set of feasible bid profiles is $X \subset X_1 \times \dots \times X_n$ and we assume that this includes the profile of null bids from each bidder. The seller has a preference ordering \succ_0 over X that guides his/her choices during the auction.

The mechanism we study is called the *ascending proxy auction*. Moves in the game are made by proxy bidders, each of whom executes an algorithm based on the preferences reported by its *client*, who may be a strategic player. Each proxy bidder begins at round one of the auction by offering the feasible bid that its client most prefers to become a winning bid. At any subsequent round in which the bid is a provisional winner, the proxy stands pat and places no new bid. At any round in which the bid is *not* a provisional winner, the proxy places the client's next-most-preferred bid. The *state* of the auction at any time is the accumulated sets of bids made on behalf of each bidder.

Let the auction round be denoted by $t = 0, 1, 2, \dots$, where round 0 is for initialization only and the actual bidding begins at round $t = 1$. Set $S_j(0) = \{\emptyset_j\}$ for each bidder j . Let $x(t)$ denote the \succ_0 -most-preferred element in $X \cap (S_1(t) \times \dots \times S_n(t))$ and, for $t \geq 1$, let $z_j(t)$ the \succ_j -most-preferred bid in $X_j - S_j(t - 1)$. Beginning with $t = 1$, the auction proceeds as follows. If $x_j(t - 1) \neq \emptyset_j$ or if $X_j - S_j(t - 1) = \emptyset$, then $S_j(t) = S_j(t - 1)$; this has the interpretation that tentatively winning bidders and bidders who have exhausted all their legal bids make no new bid. If $x_j(t - 1) = \emptyset_j$, then $S_j(t) = S_j(t - 1) \cup \{z_j(t)\}$, which has the interpretation that all bidders at the first round and tentatively losing bidders at all rounds make the new bid that is most preferred according to their clients' reported preferences.

As for the other package auctions we have studied, the seller is permitted to accept only one bid per bidder. The proxy agents bid nonstrategically, always

offering the package that would be most preferred if that bid were to win at the present round and without considering how the bid will affect the future course of the auction.

This proxy auction model does not precisely fit the rules of either of the reported experiments. The train scheduling auction allows the seller to accept multiple bids per bidder. This could be captured in the proxy auction model by allowing each shipper to be represented by a separate proxy bidder for each route. For both designs, the proxy auction omits the minimum bids set by the seller during the course of the auction, which limits the choice set from which the proxy bidders can choose, but that omission proves to be inessential.

The first achievement of the model is to shed some light on how, algorithmically, the auctions could lead to efficient outcomes.¹⁵ If the seller's preference ranking of feasible outcomes is strict, then the outcome x^* of the ascending proxy auction is a core allocation with respect to the reported preferences, regardless of what is reported. This means that the allocation is feasible and that there is no *blocking* coalition, that is, no coalition with the capacity to implement another feasible allocation that all of its members weakly prefer and some strictly prefer.

The proof is straightforward and does not depend on the assumption of transferable utility. At every round of the proxy auction, the seller selects a feasible allocation, so x^* is certainly feasible. Because no coalition of bidders can implement anything on their own without the seller, any blocking coalition must include the seller. Suppose that there is some allocation $x \neq x^*$ that the seller and some group of bidders could implement and that the bidders weakly prefer to x^* (according to their reported preferences). Then, in the proxy auction, each bidder j in the group must have offered its part x_j before the close of the auction. At the final round, x is a feasible choice for the seller, but the seller chose x^* , so the seller does not strictly prefer x to x^* . Because the seller's preferences are strict, it follows that the seller does not weakly prefer x to x^* . Thus, there is no blocking coalition, so x^* is a core allocation.

This has several important implications. One is that the outcome is efficient, which accords well with the previously cited experimental findings. A second is that in the transferable utility case, the players' payoffs are competitive, in contrast to the outcome of the Vickrey auction.

How does the proxy auction mechanism fare from an incentive perspective? The easy case for the package auctions studied earlier was the case of substitutes. To model that case, suppose first that the set of bids for each bidder consists of all possible packages and a discrete set of allowed prices, and that

¹⁵Parkes and Ungar (2002) were the first to report the efficiency of outcomes for a mechanism like the ascending proxy auction when employed on package-auction environments with quasi-linear preferences. Versions of all the theorems about the ascending proxy auction reported here appeared first in Ausubel and Milgrom (2002), with some improvements reported in Milgrom (2004).

the seller's preference maximizes his/her revenue subject to some tie-breaking rule. Bidders' preferences are similarly determined by value functions in the allowed set V .

The substitutes result is that when $V \subset V_{\text{sub}}$, truthful reporting is an ex post equilibrium, which means that no bidder could ever gain by changing his/her truthful report after having learned the reports of the other bidders. For the case of a single good, this reduces to Vickrey's observation that a dominant strategy is to bid up to one's own value in an ascending clock auction (in which the only feasible strategies are to bid up to some maximum price and then drop out). This indicates that when bid increments are small and goods are substitutes, the ascending proxy auction must approximately duplicate the results of the Vickrey auction; it leads to the unique bidder-optimal core allocation.

The ascending proxy auction differs from the Vickrey auction when goods are not substitutes. The Vickrey auction preserves the bidders' high payoffs for any reported values, even when that causes the seller's revenue to fall. The ascending proxy auction preserves the core property, sacrificing the ex post equilibrium property. That leads to the question of how badly the bidders' incentives may fail and whether the equilibrium outcome is efficient or is in the core. The second result deals with that issue.

For arbitrary valuation profiles, if π is a bidder-optimal core imputation, then the very bids $\{v_j(x_{jk}) - \pi_j\}$ that form a coalition-proof Nash equilibrium in the pay-as-bid package auction also form a Nash equilibrium in the ascending proxy auction, and the outcomes are the same. At this equilibrium, the profile of bidder payoffs coincides with the core imputation π .

For all the same reasons that I discussed in connection with the pay-as-bid auction, this equilibrium theory is less than satisfying, particularly because bidder j cannot play its equilibrium strategy without knowing π_j . Still, this theoretical analysis and others not reported here suggest that bidders may not always find it easy to improve on profit-target strategies, and that some of the good properties of equilibrium follow from those properties alone. It would seem to be worthwhile to explore the equilibrium properties more completely.

4. RECOMMENDING MECHANISMS

Package-auction theory is on the forefront of practical mechanism design. We have a pretty good choice of mechanisms to suggest when bidders regard goods as substitutes. The same is true in matching applications, in which the hard practical problems start when complementarities enter the picture, as when a hospital wants to be sure it hires an even number of doctors or when two married doctors want to find jobs in the same city (Roth and Peranson (1999)). Matching problems with substitutes are easy for much the same reasons that auctions with substitutes are easy.¹⁶

¹⁶Hatfield and Milgrom (2005).

There are other issues in practical mechanism design that we are only beginning to grapple with. Complexity is an important example. In a modestly sized package auction with $K = 30$ goods for sale, there are 2^{30} different packages, which is more than 10^9 . Even with as few as $K = 10$ goods for sale, there are more than 1,000 packages, which, except in special cases, is too many for the bidders to evaluate accurately in advance of the auction. A potential advantage of dynamical auction designs is the one familiar from the theory of planning mechanisms: they economize on communications.¹⁷ There have been some recent attempts to take that into explicit account in designing practical mechanisms.¹⁸

In terms of proven theory, the ascending proxy auction appears to be the most accomplished of the new auction mechanisms. In the complete-information case, even with complementarities, it matches the equilibrium outcomes of the pay-as-bid package auctions that are in use today. In easy environments where goods are known to be substitutes, it matches the excellent equilibrium properties of the Vickrey auction: truthful reporting is an ex post equilibrium and leads to the bidder-optimal core outcome. Unlike the Vickrey auction, there are no circumstances in which a seller in an ascending proxy auction can benefit by excluding bidders and a buyer can never benefit by bidding under multiple names. So, in terms of the theorems that we know, the auction is promising. Still, there is much that we do not know, so any conclusions must be tentative.

The incentive properties of the ascending proxy auction were derived from the observation that its outcomes are core allocations with respect to reported preferences. This suggests the possibility that other core-selecting auction mechanisms could have equally good or better properties. That is a topic for future research.¹⁹

4.1. *Package Exchanges*

Another high priority research item for mechanism designers is the analysis of *package exchanges*. Package exchanges are similar to package auctions, but they are designed to allow multiple sellers as well as multiple buyers, and sometimes to allow parties to buy some goods and sell others. Table I illustrates how the presence of multiple sellers introduces new problems. With a single seller, the Vickrey outcome with a price of 10 is a core outcome. With two sellers, as we have seen, the core is empty. If the value of 12 in the table were replaced by 24, then the core is nonempty, but a new problem emerges: how to divide the surplus among the two sellers. The rules of any package exchange therefore

¹⁷A similar point has been made in auction equilibrium models by Compte and Jehiel (2000) and by Rezende (2002).

¹⁸Parkes (2003).

¹⁹As this goes to press, some of that research has been completed. See Milgrom (2006).

involve not only an auction protocol, but also, at least implicitly, a bargaining protocol. Package exchanges are complicated objects.

One reason for the interest in package exchanges is the substantial advantages that they promise for participants. Consider, for example, a securities trader who wants to buy an option and sell related shares. The trader faces an execution risk—prices may change before the transaction is completed—and the risk is compounded when markets become less liquid. The arbitrageur could benefit from an ability to execute a package trade, so he/she buys the option only if the associated share sale is also successful.

In the United States, package exchanges are currently being studied by the FCC as a way to reallocate the spectrum currently controlled by television stations into higher valued uses. There is an active debate about whether, with modern spread-spectrum communications technologies, private licensing of spectrum even makes sense or whether some sort of regulated commons approach might use the spectrum more efficiently. For the purpose of discussing the exchange design issues, let us set aside that debate and assume that the private spectrum bands approach will be continued.

How can spectrum that is currently devoted to low-value uses like analog television broadcasting be redirected to higher value uses? One barrier is that the broadcasters' original licenses limit them to employ their spectrum for the licensed use only, rather than for general uses. One corrective policy that is sometimes suggested is simply to remove the restrictions on the broadcast licenses to allow them to migrate to higher value uses. In this case, however, such a migration involves pooling the resources of several different broadcasters, which can be difficult for just the reasons already described. There may be no core allocations or the bargaining problem may be difficult.

One idea is to try to construct a package exchange in which broadcasters who participate voluntarily would be the first to have restrictions removed from their licenses. That certainly improves participation incentives, but still leaves open questions about the rules of the exchange. Although received theories that report the impossibility of constructing an incentive-compatible exchange with a balanced budget provide some insight, they offer no useful advice about how best to proceed.

One way to approach package exchanges is as a special case of the package auctions. In the package-auction formulation, the sets of offers X_i could be specified to allow bundles with positive and negative quantities, which are interpreted as bids to buy and sell. If the auction were a spectrum exchange run by the FCC, the rules could specify that a collection of offers is acceptable only if the FCC's revenues are nonnegative. The problem with such a design is that whenever the FCC's revenues are strictly positive, some bidder has an incentive to reduce its bid to take some of the surplus.

An interesting new approach to package exchanges was developed in a recent paper by Parkes, Kalagnanam, and Eso (2002). These authors deviated from the usual tradition of mechanism design research by treating incentive

constraints differently from physical constraints. If the incentives to misreport values are not too high, then perhaps most participants will report truthfully. This attitude provides an analytical path past impossibility results like the theorems of Green and Laffont (1977) and Holmstrom (1979), which identify conditions under which dominant strategy incentives and efficient outcomes are incompatible with balanced budgets.

The new approach treats the problem of creating a direct revelation mechanism that is constrained as follows: (i) The transaction must be chosen to maximize the gains from trade, using the reported valuations. (ii) Transfers must add up to zero. (iii) Each bidder's profit, assuming its report is true, must be nonnegative. It is easy to see why a practical designer might want to impose these constraints so as to avoid subsidies and to get the participants to go along with the recommendations of the mechanism after the fact. It is well known that one cannot have all of these properties and dominant strategy incentives as well, but one can impose these constraints and then decide on payments so as to minimize the maximum incentive to deviate. The result is the *threshold exchange* design.

In the threshold exchange, assuming the reports are truthful, for each bidder j , payments are set so that the bidder's profit is $\pi_j = \max(0, \bar{\pi}_j - C)$, where $\bar{\pi}_j$ is the bidder's profit in the Vickrey mechanism. The constant C is computed so that the transfers add up to zero.

The threshold exchange is a clever design that reduces in various special cases to familiar mechanisms. For example, if a single seller reports that its cost is c and a single buyer reports that its value is b , then the mechanism calls for them to trade and to transfer $\frac{1}{2}(b + c)$ from the buyer to the seller. If two bidders report private values $b_1 > b_2$ for a good and the single seller reports cost $c < b_1$, then the mechanism calls for bidder 1 to acquire the item for a price equal to $\frac{1}{2}(b_1 + \max(b_2, c))$.

5. CONCLUSION

Package auctions, which are increasingly used in practical applications, represent a frontier in the design of practical mechanisms for resource allocation. I have tried to show that theory has an important role to play in evaluating new designs, although experiments and other methods are also indispensable. Vickrey's breakthrough mechanism, long the darling of theoretical mechanism designers, is impractical even for small auction applications, mainly because the seller's revenues can be too low, and because bidders and the seller can manipulate the design to pursue or avoid low-revenue outcomes.

There are flexible new designs, particularly the ascending proxy auction, that attractively compromise the incentive properties of the Vickrey auction with the need to avoid low revenues and that seem to correspond at least roughly to the mechanisms that are reported to have good success in economic laboratories. The experiments to date, while promising, are not yet completely

convincing. We need more experiments and more theory to understand these mechanisms better.

The package exchange problem has proved to be very difficult for practitioners. Theory suggests why this might be the case: the empty core means that attracting participants poses inevitable problems. There are, however, ways to adapt the ascending package auction to treat the package exchange problem as well as some interesting new ideas that involve rethinking the role of incentives in mechanism design.

In the [Introduction](#), I observed that the biggest surprise in the theory is a certain four-way equivalence: the very same substitutes condition that is critical for (i) the existence of competitive equilibrium is also critical for (ii) the Vickrey outcomes to lie in the core, (iii) the uniqueness of equilibrium in pay-as-bid (“menu”) auctions, and (iv) the ex post equilibrium property of the Ausubel–Milgrom ascending proxy auction. I have opted in this paper not to delve into the mathematics of the equivalence, but a sketch of the relevant mathematical ideas can be given. The condition that goods are substitutes for all bidders is a lattice-theoretic condition, equivalent to the statement that the bidders’ indirect profit functions π_j are submodular. We have used this condition to show a lattice-theoretic conclusion: the set of bidder payoffs in the core is a lattice. It follows immediately that there is a unique bidder-optimal imputation. We saw earlier that the Vickrey payoff vector gives each bidder a payoff equal to its highest payoff at any point in the core, so the Vickrey payoff is in the core if and only if there is a unique bidder-optimal imputation. The Bernheim–Whinston mechanism also selects a bidder-optimal imputation, so its multiple equilibrium outcomes vanish in this case. The connection to existence of competitive equilibrium arises because a competitive equilibrium is just a core allocation in which the goods, rather than their owners, are treated as the players. The converse theorems, which indicate that the substitutes condition is *critical* in a sense that was described precisely in the foregoing text, demonstrate that the property that the core is a lattice must be inferred from a lattice-theoretic condition, so that any preferences that are inconsistent with the substitutes condition defeat the conclusion that the bidder-optimal imputation is unique.²⁰

What I hope to have shown is that standard ideas in economic theory, particularly ideas about the core, are quite useful for the analysis and design of real-world trading mechanisms, for evaluating when the mechanisms are likely to work well, and for improving our understanding of experimental outcomes. We have used theory to identify when resource allocation problems are difficult and unorganized markets are likely to fail to suggest new designs that may

²⁰This lattice-theoretic approach also accounts for the success of other formulations. For example, the left-shoe/right-shoe formulation by Sun and Yang (2006) satisfies a lattice-theoretic substitutes condition with respect to a suitable order, in which a “larger” price vector is one that entails higher prices for left shoes and lower prices for right shoes. A similar remark applies to the recent work by Ostrovsky (2005).

help solve the difficult problems, to unify and explain the results of diverse experiments, and to generate new conjectures that can be tested in experiments. There is a palpable excitement about this area of applied mechanism design, because so much more remains to be done.

Dept. of Economics, Stanford University, Stanford, CA 94305, U.S.A.; paul@milgrom.net.

Manuscript received August, 2005; final revision received September, 2006.

REFERENCES

- AUSUBEL, L., AND P. MILGROM (2002): "Ascending Auction with Package Bidding," *Frontier of Theoretical Economics*, 1, 1019. [937,946,956,958]
- BERNHEIM, B. D., AND M. WHINSTON (1986): "Menu Auctions, Resource Allocation and Economic Influence," *Quarterly Journal of Economics*, 101, 1–31. [936,942]
- BIKHCHANDANI, S., AND J. MAMER (1997): "Competitive Equilibrium in an Exchange Economy with Indivisibilities," *Journal of Economic Theory*, 74, 385–413. [939]
- BIKHCHANDANI, S., AND J. M. OSTROY (2002): "The Package Assignment Model," *Journal of Economic Theory*, 107, 377–406. [942,946,947]
- BREWER, P., AND C. PLOTT (1996): "A Binary Conflict Ascending Price (BICAP) Mechanism for the Decentralized Allocation of the Right to Use Railroad Tracks," *International Journal of Industrial Organization*, 14, 857–886. [954]
- CANTILLON, E., AND M. PESENDORFER (2003): "Combination Bidding in Multi-Unit Auctions," Working Paper, Harvard University. [951]
- CLARKE, E. H. (1971): "Multipart Pricing of Public Goods," *Public Choice*, 11, 17–33. [942]
- COMPTE, O., AND P. JEHIEL (2000): "On the Virtues of the Ascending Auction: New Insights in the Private Value Setting," Working Paper, CERAS. [960]
- CYBERNOMICS (2000): "An Experimental Comparison of the Simultaneous Multiple Round Auction and the CRA Combinatorial Auction," available at www.fcc.gov/wtb/auctions/comin/98540191.pdf. [953]
- DE VRIES, S., AND R. VOHRA (2003): "Combinatorial Auctions: A Survey," *INFORMS Journal on Computing*, 15, 284–309. [942]
- EPSTEIN, R., L. HENRÍQUEZ, J. CATALÁN, G. Y. WEINTRAUB, AND C. MARTÍNEZ (2002): "A Combinational Auction Improves School Meals in Chile," *Interfaces*, 32, 1–14. [951]
- GREEN, J., AND J.-J. LAFFONT (1977): "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods," *Econometrica*, 45, 427–438. [945,962]
- GROVES, T. (1973): "Incentives in Teams," *Econometrica*, 61, 617–631. [942]
- GUL, F., AND E. STACCHETTI (1999): "Walrasian Equilibrium with Gross Substitutes," *Journal of Economic Theory*, 87, 95–124. [940,941]
- HATFIELD, J., AND P. MILGROM (2005): "Matching with Contracts," *American Economic Review*, 95, 913–935. [959]
- HOHNER, G., J. RICH, E. NG, G. REED, A. DAVENPORT, J. KALAGNANAM, H. S. LEE, AND C. AN (2003): "Combinatorial and Quantity Discount Procurement Auctions with Mutual Benefits at Mars, Incorporated," *Interfaces*, 33, 23–35. [951]
- HOLMSTROM, B. (1979): "Groves Schemes on Restricted Domains," *Econometrica*, 47, 1137–1144. [945,962]
- KELSO, A., AND V. CRAWFORD (1982): "Job Matching, Coalition Formation, and Gross Substitutes," *Econometrica*, 50, 1483. [937,939-941]
- KLEMPERER, P. (2000): *The Economic Theory of Auctions*. Cheltenham, U.K.: Elgar. [949]
- MILGROM, P. (2000): "Putting Auctions Theory to Work: The Simultaneous Ascending Auction," *Journal of Political Economy*, 108, 245–272. [940,941]

- (2004): *Putting Auction Theory to Work*. Cambridge, U.K.: Cambridge University Press. [941,958]
- (2006): “Incentives in Core-Selecting Auctions,” Working Paper, Stanford University. [938,960]
- MYERSON, R. B., AND M. A. SATTERTHWAIT (1983): “Efficient Mechanisms for Bilateral Trading,” *Journal of Economic Theory*, 29, 265–281. [937]
- OCKENFELS, A., AND A. E. ROTH (2003): “Last Minute Bidding and the Rules for Ending Second-Price Auctions: Evidence from eBay and Amazon on the Internet,” *American Economic Review*, 92, 1093–1103. [949]
- OSTROVSKY, M. (2005): “Stability in Supply Chain Networks,” Working Paper, Stanford University. [963]
- PARKES, D. (2003): “Auction Design with Costly Preference Elicitation,” available at www.eecs.harvard.edu/~parkes/pubs/costlyprefs.pdf. [960]
- PARKES, D., J. KALAGNANAM, AND M. ESO (2002): “Achieving Budget-Balance with Vickrey-Based Payment Schemes in Exchanges,” in *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence*. San Mateo, CA: Morgan Kaufmann, 1161–1168. [937,961]
- PARKES, D., AND L. UNGAR (2002): “Iterative Combinatorial Auctions: Theory and Practice,” in *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (AAAI-00)*. AAAI Press/MIT Press, 74–81. [942,958]
- REZENDE, L. (2002): “Mid-Auction Information Acquisition,” Working Paper, Stanford University. [960]
- ROTH, A. E., AND E. PERANSON (1999): “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design,” *American Economic Review*, 89, 748–780. [959]
- ROTHKOPF, M., T. TEISBERG, AND E. KAHN (1990): “Why Are Vickrey Auctions Rare?” *Journal of Political Economy*, 98, 94–109. [949]
- SUN, N., AND Z. YANG (2006): “Equilibria and Indivisibilities: Gross Substitutes and Complements,” *Econometrica*, 74, 1385–1402. [941,963]
- VICKREY, W. (1961): “Counterspeculation, Auctions, and Competitive Sealed Tenders,” *Journal of Finance*, 16, 8–37. [942]
- YOKOO, M., Y. SAKURAI, AND S. MATSUBARA (2004): “The Effect of False-Name Bids in Combinatorial Auctions: New Fraud in Internet Auctions,” *Games and Economic Behavior*, 46, 174–188. [948]