

Chapter 3

Ascending Proxy Auctions

Lawrence M. Ausubel and Paul Milgrom

1. Introduction

Theoretical treatments of auctions usually analyze situations in which there is some particular item to be bought or sold and the question is what auction format to use. In real auctions, however, the important planning starts much earlier and is much more encompassing. For example, in a bankruptcy auction, the auctioneer may need to decide whether to sell a whole company as a single unit or its various assets individually, what guarantees to offer concerning the conditions of its physical assets, what kinds of financing terms to require, how much time to allow buyers to obtain needed regulatory or other approvals for the acquisition, and so on.

A similarly richer set of questions arises when a buyer runs a procurement auction. For example, when the Chilean government sought to buy milk for its school milk programs, it had to decide about the sizes of the regions to be served, whether to require bonding or other performance guarantees and at what levels, how to measure and reward quality, whether and how to account for differences in past performances of suppliers, etc.

Often, these *packaging* decisions, which establish what will be bought or sold in the auction and how bids will be compared, are among the most critical decisions the auctioneer makes. Arguably, in the Chilean milk auction example, if the government were to specify that bids for service are to cover large areas of the nation, then smaller

milk suppliers would effectively be precluded from bidding. At the same time, larger bidders, enjoying economies of scale in serving a large market, could pass on part of their cost savings by reducing their bids. Given these offsetting effects, how should the service areas be specified?

In the US telecommunications spectrum auctions, sophisticated bidders anticipated the effects of packaging on the auction and lobbied the spectrum regulator for packages that served their individual interests. For example, the long-distance company MCI lobbied for a nationwide license which, it claimed, would enable cell phone companies to offer seamless coverage across the entire country. MCI knew that if such a nationwide license plan were adopted, it would exclude existing mobile telephone service providers from bidding, because those providers were ineligible to acquire new licenses covering areas that they already served. In the same proceeding, regional telephone companies such as Pacific Bell lobbied for licenses covering regional areas that fit well with their own business plans but poorly with the plans of MCI.¹

Typically, the auctioneer does not know which packaging decision is optimal. Ideally, one would like to avoid predetermining the packaging decision, instead designing an auction mechanism that allows bidders with different plans to bid accordingly. In that way, competition among the bidders would determine not just the prices but the relevant packaging decisions as well.

The extra complexity involved in package auctions has been a serious barrier to implementation and is discussed extensively in other chapters of this book. In this chapter, we set aside the issue of complexity to focus on other aspects of performance.

Many discussions of package auctions begin with what is variously called the Vickrey auction or the Vickrey-Clarke-Groves (VCG) mechanism.² As we described in chapter 1 of this volume, that mechanism and its extensions have their best performance when two conditions are satisfied: (i) the goods for sale are substitutes for all of the bidders; and (ii) bidders face no effective budget constraints. When either of these conditions fail, however, the VCG mechanism can exhibit serious performance deficiencies. We have shown that when goods are not substitutes, VCG revenues can be low or zero; revenues can decrease as more bidders are added or as some bidders raise their bids; and shill bidding and collusion can become profitable strategies for the bidders. Also, when a bidder's budget constraint is binding, the VCG auction loses its dominant-strategy property, eliminating its greatest advantage over other designs.

A second alternative is simply to solicit sealed bids and to accept the combination of bids that maximizes the seller's revenue or minimizes its cost, subject to any relevant constraints. The payment rule is "first price" (i.e., each winning bidder pays the amount of the associated bid). The Chilean milk auction cited above was run in that way, and several related applications are described in chapters 21-23. A theory of sealed-bid, first-price package auctions is developed by Bernheim and Whinston (1986). With complete information, this design has equilibria that are core allocations. However, this design has the same disadvantages as other static pay-as-bid auction formats, always forcing bidders to make guesses about the bids of others, and generally yielding inefficient outcomes in asymmetric environments with private information (including environments satisfying the substitutes condition).

In this chapter, we describe a third alternative, our ascending proxy auction, which retains some of the advantages while avoiding some of the disadvantages of the other two designs. We show that the new design duplicates the performance of the VCG mechanism when the two conditions described above hold, but leads to different results in general, avoiding the most serious shortcomings of the VCG mechanism when either condition fails. We will also show that, like the first-price package auction design, it has full-information equilibria that are core allocations.

In its simplest form, the ascending proxy auction is a direct revelation mechanism.³ Each bidder in the auction reports preferences for the contracts/packages that interest it to a proxy bidder (who may be an electronic proxy agent, or even the auctioneer), who then bids in a series of rounds on the bidder's behalf. At each round, if a given bidder is not among the provisional winners, the proxy makes whatever new bid that the bidder most prefers, according to its reported preferences. The auctioneer then considers all bids from the current and past rounds and selects his most preferred feasible collection of bids, where feasibility incorporates the constraint that the accepted bids can include at most one from each bidder. In the simplest case, the auctioneer's objective is to optimize the total price, but any other objective that leads to a unique choice is also allowed. The auctioneer's selected bids become the new provisional allocation—while the associated bidders are designated provisional winners—and the process is allowed to repeat until no new bids are submitted.

In the ascending proxy auction, both the nature of the packages or contracts between bidders and the seller and the bidders' preferences among these packages or contracts can be very general. In principle, a contract could specify a set of items, price,

quality, closing date, and so on, and any bidder preferences over such contracts could be accommodated. Although sellers will, in practice, want to limit the complexity of bids, the generality of the theory highlights the breadth of potential applications. In particular, by varying the preferences, one can accommodate bidder budget limits or even financing limits that differentially constrain what the bidder can bid for different packages.

If there is just one item for sale and only price matters, the ascending proxy auction is similar to familiar Internet auctions, where a bidder may secretly confide its values to a proxy agent (or to the auctioneer), who bids on its behalf. The one important difference is that, in the ascending proxy auction, the use of the proxy bidder is mandatory.⁴ As described in chapter 1, a single-item ascending auction with mandatory proxy bidding is strategically equivalent to the VCG mechanism for a single item (i.e., the second-price, sealed-bid auction). We show that the equivalence between the VCG mechanism and our ascending proxy auction extends to the entire range of environments in which the goods for sale are substitutes and bidders are not subject to any budget limits. This is the same as the range of environments on which the VCG mechanism avoids the problems described in Chapter 1. The most interesting question then becomes how the ascending proxy auction performs generally, when the preceding assumptions do not apply.

Our analysis has two parts. The first part characterizes the main property of the mapping from bids to outcomes: we find that the auction generally yields a core allocation with respect to the preferences reported by the bidders and the seller. The second part of our analysis studies equilibrium. We show, using an equilibrium refinement, that when preferences are quasi-linear and budget limits do not apply, the set of equilibrium allocations coincides with the set of bidder-optimal core allocations.

For additional detail about the ascending proxy auction, see Ausubel and Milgrom (2002), on which this chapter is based, or Milgrom (2004), which discusses additional applications and other related auction designs. Some aspects of the ascending proxy auction technology are further described in Ausubel and Milgrom (2001).

2. Modeling the ascending proxy auction

Denote the set of participants in the auction by $\{0, 1, \dots, L\}$, where 0 denotes the seller and $l \neq 0$ denotes a bidder. Let X_l denote the finite set of *contracts* available to bidder l . For example, in a spectrum allocation problem, X_l would contain pairs, each consisting of a set of spectrum licenses and an associated price. In that case, the relevant set X_l is finite if prices are described by positive integers and l faces some overall budget limit. For the Brewer-Plott (1996) train scheduling problem, a shipper's contract specifies a train's departure and arrival times, its direction of travel, and a price. For a procurement problem in which the eventual quantity to be purchased is uncertain, a seller's contract specifies a range of permitted quantities and an associated price schedule and may also specify contract terms concerning quality, delivery, guarantees, and so on. As described in the next section, the model even applies to public decisions with transfers, in which case each X_l is a set of pairs describing the joint decision and l 's transfer.

In the environments we have described, an *allocation* is described by a profile of contracts $x = (x_1, \dots, x_L)$. A *feasible* allocation is an element of the non-empty set $F \subset X_1 \times \dots \times X_L$. For a spectrum allocation problem, F would constrain the seller to sell each license to at most one bidder. For the Brewer-Plott train scheduling problem, F

would constrain the set of train schedules so that no trains will crash and safe spacing requirements are satisfied.

In the ascending proxy auction, *bids* are the contracts that the bidders propose. We limit attention here to modeling auctions with voluntary trade. This means that for each bidder l , the “null offer” \emptyset_l is always included in X_l . For simplicity, we assume that $(\emptyset_1, \dots, \emptyset_L) \in F$. This means that the seller is not compelled to sell anything to anyone, which ensures that there is always at least one feasible allocation for the seller.

Suppose that each bidder l ranks any allocations x based only on its own contract x_l and let \succ_l denote a strict preference ordering over the finite set X_l .^{5,6} In the proxy auction, the bidder reports such a preference ordering to its proxy bidder. The auctioneer specifies a preference ordering \succ_0 over the set of feasible bid profiles F . Given any set of bidders S , denote the set of feasible bid profiles for the coalition consisting of those bidders and the seller by $F_S = \{x \in F \mid x_l = \emptyset_l \text{ for all } l \notin S\}$.

Once the bidders report their preferences, the *ascending proxy auction* operates in a series of rounds, where the state variable for the auction is the profile $(B_l^t)_{l=1, \dots, L}$ of bid sets and the provisional allocation x_0^t . The seller initializes the bid sets by entering the null bid on each bidder’s behalf: $B_l^0 = \{\emptyset_l\}$ and initializes the provisional allocation x_0^0 by assigning each bidder its null contract: $x_{0l}^0 = \emptyset_l$. Each round t of the ascending proxy auction proceeds as described below, in which the notation “max” applied to any choice set is defined with respect to the chooser’s reported preference ordering.

1. For each bidder, determine the set of *available new bids*, A_l^t , which comprises the individually rational bids not yet made: $A_l^t = X_l - \{x_l \succ_l \emptyset\} - B_l^{t-1}$.
2. Any bidder l who is not a provisional winner and who has not exhausted its profitable bids offers its \succ_l -most preferred contract in A_l^t . Formally, $[x_{0l}^t = \emptyset, A_l^t \neq \emptyset] \Rightarrow B_l^t = B_l^{t-1} \cup \{\max A_l^t\}$.
3. All other bidders make no new bids: $[x_{0l}^t \neq \emptyset \text{ or } A_l^t = \emptyset] \Rightarrow B_l^t = B_l^{t-1}$.
4. If there are no new bids at round t ($B^t = B^{t-1}$), then the auction terminates and the provisional allocation x_0^t becomes the final allocation. Otherwise ($B^t \neq B^{t-1}$), the auctioneer determines its current feasible set, $F^t = F \cap \bigcap_{l=1}^L \{x \mid x_l \in B_l^t\}$, selects its \succ_0 -most-preferred element $x_0^{t+1} = \max F^t$, and iterates with round $t + 1$.

The ascending proxy algorithm determines a mapping from reported preferences to allocations. The first step in the analysis is to establish a key property of that mapping, namely, that it picks points in the core.

Theorem 1. The allocation determined by the ascending proxy auction is a stable allocation (and hence an NTU-core allocation) with respect to the reported preferences.⁷

Proof. Let T be the final round and let x_0^T be the allocation determined by the ascending proxy auction. By construction, x_0^T is individually rational for bidders ($x_l^T = \emptyset$ or $x_l^T \succ_l \emptyset$) and since only the null allocation is feasible for coalitions excluding the seller, no coalition of bidders alone can block x_0^T . It therefore suffices to show that no coalition of the seller and one or more bidders can block.

Let S be any non-empty set of bidders. Suppose $x \in F_S$, so that for all $l \notin S$, $x_l = \emptyset_l$. Contrary to our conclusion, suppose that the coalition consisting of the seller and the bidders in S blocks x_0^T with x , that is, $x \succ_0 x_0^T$ and for all $l \in S$, $x_l = x_{0l}^T$ or $x_l \succ_l x_{0l}^T$. Consider the components of the auction outcome x_0^T . By construction, whether l is a losing bidder ($x_{0l}^T = \emptyset_l$) or a winning bidder ($x_{0l}^T \neq \emptyset_l$), by the end of the auction, the proxy for l in the auction has made every bid that l weakly prefers to x_{0l}^T . In particular, the proxy has bid x_l , so $x_l \in B_l^T$. Hence, $x \in F^T = \bigcap_{l=1, \dots, L} \{x \mid x_l \in B_l^T\} \cap F$. Then, by step 4 of the algorithm, it cannot be that $x \succ_0 x_0^T$, contradicting the hypothesis that x_0^T is blocked. ■

Theorem 1 provides an important hint about how the ascending proxy auction is related to the Vickrey auction. In chapter 1 of this volume, we showed that the Vickrey auction leads generally to core allocations only if the goods for sale are substitutes and bidder budgets are unlimited. Failure of the outcome to lie in the core, we showed, is associated with low seller revenues, non-monotonic prices (adding bidders or increasing bidder values may reduce prices), existence of profitable shill-bidding opportunities by individual bidders,⁸ and existence of profitable joint deviations by groups of losing bidders.

Theorem 1 and its proof are closely related to similar results and proofs about deferred acceptance algorithms in matching theory. (For example, see Gale and Shapley 1962, Kelso and Crawford 1982, and Roth and Sotomayor 1990.) In all of these algorithms, one side of the market makes a sequence of bids and continues to add bids when its old bids are rejected. Further connecting the results is Hatfield and Milgrom's

(2004) finding that, if the final allocation from such a *cumulative offer process* is feasible, then it is a stable allocation.

3. *The transferable utility case*

Theorem 1’s conclusion that the outcome is a stable allocation (and hence an NTU-core allocation) with respect to reported preferences has three important limitations. First, it does not identify which NTU-core outcome among the possibly many such outcomes is selected. Second, it is silent about whether or when the reported preferences are likely to coincide with or even resemble the bidders’ actual preferences, which importantly affects the interpretation of the result. Third, it is silent about how the outcome relates to the NTU-core when bidders play an equilibrium strategy that misreports their actual preferences. To answer the associated questions, we specialize the model and introduce additional assumptions.

Our first change is to focus attention on the case in which the auction is conducted by a seller who ranks feasible collections of bids according to their total price. In this setting, an allocation is a pair (z, p) where $z = z_0 = (z_1, \dots, z_L)$ is a *decision* and p is a vector of *cash transfers*. In a typical package auction, the decision describes which bidder gets which goods, but there are other possibilities as well. For example, in a club or social organization, z may describe the levels of investment in club facilities and amenities and the ascending proxy auction may be used to decide both z and which bidders will be members. In this setting, it is convenient to apply the feasibility criterion to the decision z , rather than to the allocation (z, p) . The condition in the club version of the problem that all “members” enjoy the same services \hat{z} is $F \subset \bigcup_{\hat{z} \in \hat{Z}} \{(z_1, \dots, z_L) \mid (\forall l) z_l \in \{\hat{z}, \emptyset_l\}\}$.

We assume that the bidders' valuations are quasi-linear, so a bidder's payoff can be written as $v_l(z) - p_l$ or $v_l(z_l) - p_l$. In this *transferable utility* case, the total payoff is $\sum_{l=1}^L v_l(z_l)$ and variations in the payment profile p transfer utility among the participants.

Finally, to reinforce the idea that there is just one seller, we assume that if $z = (z_1, \dots, z_L) \in F$ is a feasible decision, then so is (\emptyset_l, z_{-l}) . This ensures that one can exclude bidder l and its contract without affecting the feasibility of the sale, as is the case in many package auctions.

The ascending proxy auction modeled in the preceding section uses finite feasible sets for the bidders. In the transferable utility case, it simplifies the analysis to imagine that the monetary unit is very small, so that the NTU-core is closely approximated by the traditional transferable utility (TU) core. We will make our arguments below as if this approximation were exact. Also, since the theory was based on the assumption that seller preferences are strict, we need a method to resolve ties. Many different procedures are satisfactory. For example, it suffices to imagine that the seller breaks ties by adding a negligibly small, random, negative amount to each bid.

To describe the TU core, one must first define the coalitional game (L, w) that is associated with the trading model. The set of players is just as above: $L = \{0, \dots, L\}$. The coalitional value function is defined as follows:

$$w(S) = \begin{cases} \max_{x \in X} \sum_{l \in S} v_l(x_l), & \text{if } 0 \in S, \\ 0, & \text{if } 0 \notin S. \end{cases} \quad (1)$$

The value of a coalition is the maximum total value the players can create by trading among themselves. If the seller is not included in the coalition, that value is zero.

In TU-games, it is traditional to define the core by suppressing the allocation and focusing on the imputed payoffs, or *imputations*, associated with any feasible allocation. The TU-core in this case is a set of payoffs, as follows:

$$\text{Core}(\mathbf{L}, w) = \left\{ \pi : w(\mathbf{L}) = \sum_{l \in \mathbf{L}} \pi_l, w(S) \leq \sum_{l \in S} \pi_l \text{ for all } S \subset \mathbf{L} \right\}. \quad (2)$$

As new bids are introduced in the ascending proxy auction, the bidders move down their preference lists, so the payoff associated with each new bid is less than that associated with all earlier bids by the same bidder. It is convenient to track the progress of the auction in terms of the payoff associated with each bidder's most recent bid. Let π_l^t be that payoff for bidder l at round t . The highest bid that l has made for decision z_l by round t is within one bid increment of $\max(0, v_l(z_l) - \pi_l^t)$. In our calculations below, we treat this estimate of the bid as exact.

Imagine that the seller regards the auction as serving two purposes: creating a coalition to utilize its goods and the bidders' additional resources and expertise; and determining how to distribute the value created by the winning coalition. From that perspective, at every round of the auction, there is an implicit bid by every coalition of bidders for the seller's goods. Given our previous formula describing the bids, the maximum total price offered at round t by coalition S is

$\max_{z \in F} \sum_{l \in S-0} \max(0, v_l(z_l) - \pi_l^t) = \sum_{l \in S-0} \max(0, v_l(z_{Sl}^*) - \pi_l^t)$. Notice that if coalition S includes any bidder n for whom $v_n(z_{Sn}^*) - \pi_n^t < 0$, then the coalition $S - n$ bids as much as coalition S by using the same allocation among its continuing members (and setting $z_{S-n,n}^* = \emptyset_n$). If π_0^t is the highest total bid at round t , then

$$\max_S \left(w(S) - \sum_{l \in S-0} \pi_l^t \right) = \pi_0^t, \quad (3)$$

so for all S , $w(S) \leq \sum_{l \in S} \pi_l^t$. This means that in the transferable utility game with coalitional value function w , no coalition can block the payoff vector π^t . At the last round T , the auction payoff vector π^T is the one associated with the actual final allocation, so it is also feasible and hence a core payoff imputation.

This analysis suggests an interesting interpretation of the proxy auction for the transferable utility case. For concreteness, think of the bidders in the auction as workers seeking employment and the seller in the auction as a supplier of capital assets. Imagine that there is a large set of competing brokers who may hire workers and buy assets from the seller to run a business. Each broker considers only one combination of workers whom it might employ. At round t , a typical bidder l demands a payoff or “wage” of π_l^t and a broker representing coalition S is driven by competition with identical brokers to offer $w(S) - \sum_{l \in S-0} \pi_l^t$ to acquire the seller’s capital. At each round, there is a provisional winning coalition S and workers not in that coalition are unemployed. Any unemployed workers who have demanded positive wages reduce their wage demands by one unit, and the brokers then bid again. The process continues iteratively; it ends when a competitive equilibrium wage vector is found.

In this interpretation, the auction is simply a *tatonnement* device to identify a competitive equilibrium. A (TU-)core imputation is a vector of competitive equilibrium prices specifying a wage for each worker and a price for the seller’s bundle of assets. The (TU-)core conditions state that the winning broker just breaks even and that there exist no profit opportunities for any broker at the given prices.

With transferable utility, we have the following special case of Theorem 1:

Theorem 2. In the transferable utility model, the payoff imputation determined by the ascending proxy auction is a core imputation with respect to the reported preferences.

4. Profit-target strategies and equilibrium

We turn next to our analysis of bidder incentives and equilibrium. Our first result shows that one can sometimes limit attention to the class of π_l -profit-target or *semi-sincere* strategies, in which a bidder reports its value for each decision z to be $\max(0, v_l(z) - \pi_l)$.⁹ That is, l reduces all of its reported values by essentially the same fixed amount π_l , which may be regarded as its minimum profit target. We say “essentially” because negative values are reported to be zero.

Theorem 3. In the transferable utility model, given any pure strategy profile for the other bidders, bidder n has a best reply that is a profit-target strategy.

Proof. Fix bidder n and any pure strategy report profile \hat{v}_{-n} for the competing bidders. Define $\Pi_n(\hat{v})$ to be n 's profit for any report profile \hat{v} . Let $\pi_n = \max_{\hat{v}_n} \Pi_n(\hat{v})$ be n 's maximum profit, $v_n^* \in \arg \max_{\hat{v}_n} \Pi_n(\hat{v})$ be any best reply, and z^* be the corresponding decision selected by the ascending proxy auction. Since the outcome is in the core, for all decisions z , $\sum_{l=1}^L \hat{v}_l(z^*) \geq \sum_{l=1}^L \hat{v}_l(z)$. Also, since n pays $v_n(z^*) - \pi_n$, we have $\hat{v}_n(z^*) \geq v_n(z^*) - \pi_n$.

Consider n 's π_n -profit-target reporting strategy: $\tilde{v}_n(z) = \max(0, v_n(z) - \pi_n)$. Since bids are non-negative, we have that for all z , $\tilde{v}_n(z^*) + \sum_{l \neq n} \hat{v}_l(z^*) \geq \sum_{l \neq n} \hat{v}_l(z)$. So, the

decision z selected when n reports is one for which $0 \leq v_n(z) - \pi_n$, that is, n is a winning bidder and pays a price of at most $\tilde{v}_n(z)$. It follows that its profit is at least

$v_n(z) - \tilde{v}_n(z) = \pi_n$, so the π_n -profit-target reporting strategy is also a best reply. ■

Profit-target strategies are of interest for several reasons. First, if all bidders adopt such strategies and if the selected decision is not one for which any losing bidder l has set $\pi_l = 0$, then the decision maximizes total values, that is, it is efficient. Second, a bidder who uses such a strategy bids its full incremental value for changing from any decision z to any alternative decision z' , which means it cannot be part of any collusive, price reducing agreement. For example, suppose the decision concerns the allocation of two licenses. Suppose, given the bids, the first license is won by bidder 1 and the second by bidder 2. If bidder 1 adopts any profit target strategy, then bidder 2 must pay at least 1's incremental value, $v_1(12) - v_1(1)$, to win license 2. The conclusion in theorem 3 that each bidder always has a best reply that is a profit-target strategy means that there does not exist any profile of reports for the other bidders that can deter the bidder from bidding aggressively, offering to pay up to their full incremental values for different or extra licenses. More generally, within the rules of the proxy auction, there is no strategy profile that can deter a bidder from bidding to change the predicted decision to one that it prefers.

The next theorem identifies full information Nash equilibria of the ascending proxy auction in profit-target strategies. The equilibrium outcome will be a *bidder-optimal point in the core*, meaning that the payoff profile $\pi \in \text{Core}(\mathbf{L}, w)$ and that there is no $\pi' \in \text{Core}(\mathbf{L}, w)$ with $\pi' \neq \pi$ and $\pi'_l \geq \pi_l$ for every bidder l . Notice that the outcome is a

bidder-optimal core allocation with respect to the bidders' *actual* preferences. In particular, the outcome is efficient and has the desirable revenue properties associated with core allocations. This duplicates a property previously identified by Bernheim and Whinston (1986) for sealed-bid, first-price package auctions.

Theorem 4. In the transferable utility model, for every bidder-optimal point π in the core, the strategy profile in which each bidder l plays its π_l -profit target strategy is a Nash equilibrium with associated profit vector π . Moreover, if \tilde{v} is a Nash equilibrium in semi-sincere strategies at which losing bidders bid sincerely, then $\pi(\tilde{v})$ is a bidder-optimal point in $Core(\mathbf{L}, w)$.

Proof. Suppose that the proposed strategy profile is not an equilibrium. Then there is some player l and some unilateral deviation for that player that leads to a winning coalition T ($T \ni l$) and profit outcome vector $\hat{\pi}$. For bidder l , $\hat{\pi}_l > \pi_l$, and for all bidders $k \in T$, the proxy strategies imply that $\hat{\pi}_k \geq \pi_k$.

Since π is bidder-optimal, there is a coalition S such that $l \notin S$ and $w(S) = \sum_{k \in S} \pi_k$ (Otherwise, for some $\varepsilon > 0$, there would be a point in the core at which l gets $\pi_l + \varepsilon$, the seller gets $\pi_0 - \varepsilon$, and others payoffs are as specified by π , contradicting bidder-optimality.) Let $\beta(S)$ and $\beta(T)$ denote the highest total revenue associated with bids by the bidders in coalitions S and T during the proxy auction, given the specified deviation by bidder l . We show that $\beta(S) > \beta(T)$, contradicting the hypothesis that T is the winning coalition. Indeed, using (3):

$$\begin{aligned}
\beta(S) &\geq w(S) - \sum_{k \in S-0} \max(\pi_k, \hat{\pi}_k) \\
&> w(S) - \sum_{k \in S-0} \pi_k - \sum_{k \in T-0} \max(0, \hat{\pi}_k - \pi_k) \\
&= \pi_0 - \sum_{k \in T-0} \max(0, \hat{\pi}_k - \pi_k) \\
&\geq w(T) - \sum_{k \in T-0} \pi_k - \sum_{k \in T-0} \max(0, \hat{\pi}_k - \pi_k) \\
&= w(T) - \sum_{k \in T-0} \hat{\pi}_k \\
&= \beta(T)
\end{aligned} \tag{4}$$

The first step in (4) follows from the proxy rules: any losing bidders in S stop bidding only when their potential profits reach the specified levels. The strict inequality in the second step follows because $l \in T \setminus S$ and $\hat{\pi}_l > \pi_l$. The third step follows by selection of S , the fourth because $\pi \in \text{Core}(\mathbf{L}, w)$, and the fifth and sixth by the definitions of T , $\hat{\pi}$ and $\beta(T)$.

If π is not bidder-optimal in the core, then there exists l and $\pi'_l > \pi_l$ such that $(\pi'_l, \pi_{-l}) \in \text{Core}(\mathbf{L}, w)$. Bidder l can deviate by reporting $\max\{\tilde{v}_l - \pi_l + \pi'_l, 0\}$ instead of reporting \tilde{v}_l , thereby increasing its profits to π'_l . ■

The examples in chapter 1 showed that Vickrey auction outcomes can fail to be in the core because seller revenues are too low. That conclusion, and the connection to matching theory, may seem to suggest that bidder-optimal outcomes are seller-pessimal ones, that is, outcomes that minimize the seller's revenues among all core allocations.

Here is a counter-example to refute that suggestion. There are three items for sale, denoted by A, B and C. Values for various packages are described in the following table:

Bidder	A	B	C	AC	BC
1	2				
2		2			
3			10		
4				8	8

We identify payoffs by a 5-vector with the payoff of the seller (player 0) listed first. Using the tabulated values, let us verify that $X=(8,2,2,2,0)$ and $Y=(10,0,0,4,0)$ are bidder-optimal core allocations. Observe, first, that X and Y are individually rational and feasible, with a total payoff of 14.

Any blocking coalition must involve the seller and pay him more than 8, so we may limit attention to coalitions with value strictly exceeding 8. We may further limit attention to minimal coalitions with any particular value.

There are six coalitions that meet these criteria. Coalition 0123 is the unique minimal coalition with value 14 and it receives 14 at both of the identified imputations, so it does not block. Coalitions 013 and 023 have value 12 and receive 12 or 14 at the two imputations, so they do not block. Coalitions 014 and 024 have value 10 and receive 12 or 14 at the two imputations, while coalition 03 has value 10 and receives 10 or 14 at the two imputations, so none of those coalitions block. We conclude that the imputations X and Y are unblocked and hence are core imputations.

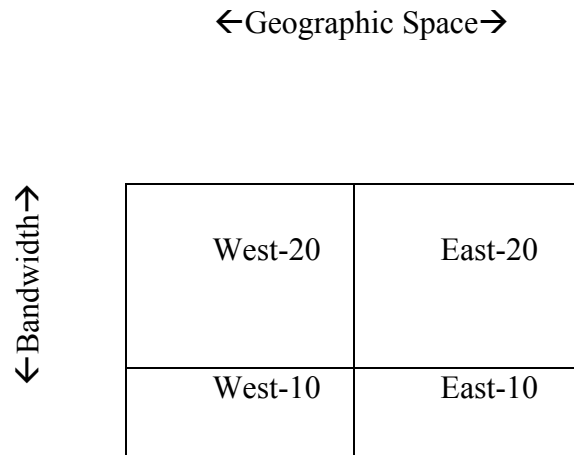
Imputation X , in which the seller gets 8, is seller pessimal, because: (i) bidder 4 must get zero at every core imputation; and (ii) the coalition 04 must get at least 8. Hence, X is a bidder-optimal core imputation. At imputation Y , bidder 3 gets his Vickrey payoff of 4, which, as we found in chapter 1, is his best core payoff. So, any bidder-preferred core imputation must have the form $(10-x-y, x, y, 4, 0)$. If $x > 0$, this imputation is blocked by coalition 024 and if $y > 0$, it is blocked by coalition 014. Since any bidder-preferred imputation is blocked, Y is a bidder-optimal core imputation.

This example identifies two-bidder optimal core allocations with different values for the seller, so bidder optimality is not the same as seller pessimality.

5. When sincere bidding is an equilibrium

We found in Chapter 1 that, if the Vickrey payoff vector is an element of the core, then there is a unique bidder-optimal point in the core, which eliminates a fundamental bargaining problem among the bidders. It might then seem that, under this condition, each bidder could simply use the straightforward bidding strategy and this would yield an equilibrium corresponding to the unique bidder-optimal point. However, this conclusion is not quite right, as the following example demonstrates.

Suppose that there are four spectrum licenses. In order to understand that the following bidder valuations are sensible, it is helpful to depict the licenses as follows:



There are five bidders. Bidder 1 desires a 10-MHz band of spectrum covering both East and West. Bidders 2 and 3 desire a 20-MHz band of spectrum covering both East and West. Bidder 4 wants the full 30-MHz of spectrum in the East. Bidder 5 wants the full 30-MHz of spectrum in the West. Thus:

$$v_1(\text{West-10, East-10}) = 10,$$

$$v_2(\text{West-20, East-20}) = 20,$$

$$v_3(\text{West-20, East-20}) = 25,$$

$$v_4(\text{East-20, East-10}) = 10, \text{ and}$$

$$v_5(\text{West-20, West-10}) = 10,$$

with all singletons and all other doubletons valued at zero.

Observe that the Vickrey payoff vector, $(20, 10, 0, 5, 0, 0)$, is an element of the core, corresponding to Bidder 3 paying 20 for his desired licenses and Bidder 1 paying 0 for his desired licenses. Nevertheless, straightforward bidding is likely to lead Bidder 1 to pay a positive price.¹⁰ In the ascending proxy auction, bidder 1 could do better by setting profit target at or slightly below its Vickrey profit, ensuring that it never pays a substantial positive price for its licenses.

5.1 Bidder-submodular values

In this subsection and the next, we explore conditions under which there is a truthful reporting equilibrium of the ascending proxy auction. In this subsection, we investigate a restriction on the coalitional value function which implies that truthful reporting is an equilibrium. In the next subsection, we explore a corresponding condition on valuation of goods.

In Chapter 1, we defined the coalitional value function w to be *bidder-submodular* if bidders are more valuable when added to smaller coalitions than when added to larger coalitions. Formally, the condition is that for all $l \in L - 0$ and all coalitions S and S' satisfying $0 \in S \subset S'$, $w(S \cup \{l\}) - w(S) \geq w(S' \cup \{l\}) - w(S')$.¹¹

Since Shapley (1962), bidder-submodularity has been called the condition that “bidders are substitutes,” but that description does not coincide with the usual economic meaning of substitutes. For suppose there were a labor market in which the bidders could be hired and a seller who hired bidders to form coalition S earned a profit of $w(S)$ minus the total wages paid. The bidder-submodularity condition is necessary, but not sufficient, for bidders to be substitutes in that labor market. For example, with three bidders, suppose that $w(02) = w(03) = 1$, $w(01) = w(023) = w(012) = w(013) = w(0123) = 2$, and all other values are zero. By inspection, w is bidder-submodular. Bidders, however, are not substitutes, because changing the wage profile from $(1.7, 0.8, 0.8)$ to $(1.7, 1, 0.8)$ changes the hiring decision from $(0, 1, 1)$ to $(1, 0, 0)$. That is, *increasing* the wage of bidder 2 *reduces* the demand for bidder 3, contrary to the standard economic definition of substitutes.

In Theorem 6 of chapter 1, we established that the following three statements are equivalent:

- (1) The coalitional value function w is bidder-submodular.
- (2) For every coalition S that includes the seller, the restricted Vickrey payoff vectors all lie in the cores of the corresponding restricted games:

$$\bar{\pi}(S) \in Core(S, w).$$

- (3) For every coalition S that includes the seller, there is a unique core point that is unanimously preferred by the buyers and, indeed:

$$Core(S, w) = \{\pi_S \mid \sum_{l \in S} \pi_l = w(S), 0 \leq \pi_l \leq \bar{\pi}_l(S) \text{ for all } l \in S - 0\}. \quad (5)$$

Here we establish an additional result about bidder-submodular values.

Theorem 5. Suppose that the coalitional value function is bidder-submodular. Then, truthful reporting is a Nash equilibrium strategy profile of the ascending proxy auction and leads to the Vickrey outcome: $\pi^T = \bar{\pi}$.

Proof. We first establish that truthful reporting leads to the Vickrey payoff vector. Suppose there is some round t at which $\pi_l^t < \bar{\pi}_l$. We show that l is necessarily part of the winning coalition at that round. Let S be any coalition including the seller but not bidder l . Then,

$$\begin{aligned}
w(S) - \sum_{k \in S} \pi_k^t &< w(S) - \sum_{k \in S} \pi_k^t + (\bar{\pi}_l - \pi_l^t) \\
&= w(S) - \sum_{k \in S \cup \{l\}} \pi_k^t + w(L) - w(L - l) \\
&\leq w(S) - \sum_{k \in S \cup \{l\}} \pi_k^t + w(S \cup l) - w(S) \\
&= w(S \cup \{l\}) - \sum_{k \in S \cup \{l\}} \pi_k^t
\end{aligned} \tag{6}$$

So, l 's profit π_l^t is at least $\bar{\pi}_l$ minus one bid increment ε . Taking the bid increment to zero for the ascending proxy auction proves that for $l \neq 0$, $\pi_l^T \geq \bar{\pi}_l$, and the reverse inequality follows from theorem 4 of Chapter 1.

Second, we show that truthful reporting is a best response to all other bidders reporting truthfully. For any bidder l and any report by that bidder, theorem 4 of Chapter 1 implies that the payoff to coalition $L - l$ is at least $w(L - l)$. Since the total payoff to all players is at most $w(L)$, l 's payoff to any strategy is bounded above by $\bar{\pi}_l = w(L) - w(L - l)$, which is the payoff that results from truthful reporting. ■

5.2 When goods are substitutes

In the preceding subsection, we studied bidder-submodularity, which is a restriction on the coalition value function. In auction models, however, coalition values are not

primitive—they are derived from individual package values. It is therefore natural to ask: what conditions on bidder valuations *imply* bidder-submodularity of the coalitional value function?

A key to the answer lies in a characterization of bidder-submodularity that we developed earlier. In Theorem 9 of chapter 1, we established that if there are at least four bidders and if the set of possible bidder values \mathcal{V} includes all additive values, then the following three conditions are equivalent:

- (1) \mathcal{V} includes only valuations for which goods are substitutes.
- (2) For every profile of bidder value functions drawn for each bidder from \mathcal{V} , the coalitional value function is bidder-submodular.
- (3) For every profile of bidder value functions drawn for each bidder from \mathcal{V} , $\bar{\pi} \in \text{Core}(L, w)$.

According to this theorem, if goods are substitutes for each bidder, then the coalitional value function is bidder-submodular. Consequently, Theorem 5 can be recast in terms of substitutes preferences, as follows:

Theorem 6. If goods are substitutes for all bidders, then truthful reporting is a Nash equilibrium strategy profile of the ascending proxy auction and leads to the generalized Vickrey outcome: $\pi^T = \bar{\pi}$. Moreover, if bidders are restricted to report preferences such that goods are substitutes, and if bidder l 's actual preferences have the property that goods are substitutes, then it is a *dominant strategy* for bidder l to report truthfully.

The theorem also has a converse. If we include any goods values for which goods are not substitutes, then we cannot ensure that the coalition value function is bidder-

submodular, so the theory of the preceding subsection does not apply. What is more, we also cannot ensure that the Vickrey outcome is a core outcome, so there cannot be an incentive for bidders to report truthfully in the ascending proxy auction. The reason is that, as observed in chapter 1, the Vickrey auction is the unique auction that, on a wide class, has the dominant-strategy property, leads to efficient outcomes, and takes only a zero payment from losing bidders.

When goods are not substitutes, the theorem implies that the ascending proxy auction necessarily departs from the VCG results, because it selects core allocations with respect to reported preferences, while the VCG mechanism does not. Unlike the VCG mechanism, the proxy auction does not have a dominant-strategy solution in every private-values environment. However, the proxy auction has some offsetting advantages, at least at its full-information equilibrium.

6. Comparisons of the Vickrey and ascending proxy auctions

In Chapter 1, we found that the failure of the substitutes condition and of bidder-submodularity was closely connected to some extreme possibilities for manipulation in the Vickrey auction, including shill bidding and loser collusion. We also established that these possibilities for manipulation were intimately related to a failure of monotonicity of revenues in the set of bidders.

These shortcomings of the Vickrey auction contrast sharply with the properties of the ascending proxy auction. For the latter, to check bidder monotonicity, we need to identify the “equilibrium revenues.” We focus attention on the equilibrium of the proxy auction that is consistent with the selection in theorem 4 and that minimizes revenues among those. Using the characterization of the theorem, the minimum equilibrium

revenue is $\min \pi_0$ subject to $\sum_{l \in S} \pi_l \geq v(S)$ for every coalition $S \subset L$. It is straightforward to see that introducing additional bidders or increasing the reports of existing bidders must (weakly) increase π_0 in this formulation. Suppose that $\hat{v}_l(\bullet) \geq v_l(\bullet)$ for all $l \in L - 0$. Consider any profit allocation $(\hat{\pi}_l)_{l \in L}$ that satisfies $\sum_{l \in S} \hat{\pi}_l \geq \hat{v}(S)$ for all $S \subset L$. Observe that $\sum_{l \in S} \hat{\pi}_l \geq v(S)$ is also satisfied for all $S \subset L$, implying that the minimum π_0 (subject to $\sum_{l \in S} \pi_l \geq v(S)$) is weakly lower than the minimum $\hat{\pi}_0$ (subject to $\sum_{l \in S} \hat{\pi}_l \geq \hat{v}(S)$). In this sense, the ascending proxy auction satisfies bidder monotonicity.

One weakness of the Vickrey auction is its susceptibility to collusion, even by coalitions of losing bidders. The ascending proxy auction is clearly immune to loser collusion. Let S be any coalition of losing bidders. Then the prices paid by the complementary set, $L \setminus 0 \setminus S$, of bidders in the ascending proxy auction sum to at least $w(S)$; otherwise, coalition S would outbid the winners. So, to become winning bidders, members of the coalition S would have to raise their total winning bid above $w(S)$, suffering a loss in the process. Hence, a coalition consisting only of losing bidders has no profitable deviation.

A more severe failing of the Vickrey auction is its susceptibility to shill bidders, that is, a bidder can sometimes increase its payoff by employing two or agents to bid on its behalf. The ascending proxy auction is immune to the use of shills in the following sense: Given any pure strategy profile for the bidders besides l , there is a best reply for bidder l that does not use shill bids. In particular, this implies that at any pure strategy

Nash equilibrium, there is no gain to any deviation using shill bids, in contrast with the result for Vickrey auctions.

Since this result can be proved using the same reasoning as for Theorem 3, we simply sketch the argument here. Bidder l , or its several agents acting together, can acquire a set of goods A in the ascending proxy auction only if its final bid for that set exceeds the incremental value of A to the coalition $L-l$. This minimum price for A can only be increased by adding losing bidders to the coalition $L-l$, so bidder l cannot reduce its minimum price for any bundle it might acquire by using shills. Moreover, bidder l can win the bundle A at the minimum price by bidding just for A and offering the incremental value. Therefore, bidder l can achieve its maximum payoff using a strategy that does not employ shills.

Each of the above comparisons has been made using one of two assumptions: either goods are substitutes or bidders have complete information about values. Although there are no theorems to identify how far our conclusions extend to environments where both conditions fail, there are robust examples showing that there is a class of environments with incomplete information and without the substitutes property that still have properties similar to those that we have described.

To illustrate, consider the following incomplete information, private-values extension of an example from Chapter 1. There are two items and two bidders. Bidder 1 wants only the package of two licenses and values it at v_1 , which is a random variable with support $[0,3]$. Bidder 2 values each license singly at v_2 and values the package of two licenses at $2v_2$, where v_2 is a random variable with support $[0.3,0.7]$. Each bidder i

knows the realization of v_i but only the distribution of v_j ($j \neq i$). In the Vickrey auction, if bidder 1 is expected to report its value truthfully, then bidder 2 would benefit by participating in the auction under two names—say, “bidder 2” and “bidder 3”—each of whom bids \$3 for a single license. The result would be that bidder 2 receives both licenses for a price of zero, which is the same conclusion as in the complete-information case.

By contrast, in the proxy auction, it is optimal for bidder 2 to bid straightforwardly, that is, to bid v_2 for each single license and to bid $2v_2$ for the package of two licenses. If bidder 2 instead bids using two identities, his optimal strategy is still to place bids that sum to $2v_2$ for the package of two licenses. In either event, the result is that bidder 2 receives both licenses if and only if $2v_2 \geq v_1$, and then for a price of v_1 . In this example with private information and without substitutes preferences for bidder 1, the Vickrey auction remains vulnerable to shill bidding, while the ascending proxy auction implements the Vickrey outcome without that vulnerability.

7. Conclusion

Our analysis of the ascending package auction may explain the efficient outcomes that sometimes emerge in experiments with package bidding and entail new predictions as well. If bidders bid “straightforwardly,” the auction outcome is not only efficient, but also a core allocation of the exchange game.

The ascending proxy auction is a new kind of deferred acceptance algorithm, related to the algorithms studied in matching theory. Unlike the sealed-bid, first-price package auction, this new design duplicates the advantages of the Vickrey auction in

environments where goods are substitutes, but the new auction design also has significant advantages compared to the Vickrey auction.

First, the ascending proxy auction avoids the very low revenues possible in the Vickrey auction, because the former always selects core allocations both when bidders bid straightforwardly and at their full information equilibria. By contrast, the Vickrey auction is assured to select core allocations only when goods are substitutes.

Second, the ascending proxy auction also avoids much of the vulnerability of the Vickrey auction to shill bids and to collusion by coalitions of losing bidders. Whereas a bidder in a Vickrey auction may sometimes find it profitable to bid under multiple names, there exists no pure strategy profile in the ascending proxy auction that ever makes such a behavior advantageous. And, whereas losing bidders in a Vickrey auction can sometimes collude profitably, that is never possible at the identified equilibria of the ascending proxy auction.

Third, the ascending proxy auction can be implemented as a multi-stage process, in which bidders first specify initial proxy values and later have one or more opportunities to revise their proxy bids. That implementation, like other sequential communication protocols, economizes on the amount of information that bidders need to communicate. It may also economize on bidders' costs of evaluating packages by allowing them to focus their efforts on the packages that they have a reasonable chance to win based on the bids made by competitors earlier in the auction.

Besides the basic ascending proxy auction, we have also introduced a generalized version that applies for much more general preferences and constraints than are permitted in the standard Vickrey model. For example, it applies to models in which bidders are

budget-constrained and to procurement problems with their characteristic complex constraints and multi-faceted selection criteria. With complex auctioneer preferences, the Vickrey auction may not even apply, and with bidder budget constraints it may lead to inefficient outcomes, but the generalized ascending proxy auction selects allocations in the NTU-core. We believe that this family of auction designs holds considerable promise for a variety of practical applications.

References

- Ausubel, Lawrence M. (1997), "On Generalizing the English Auction," unpublished working paper, <http://www.ausubel.com/auction-papers/97wp-generalizing-the-english-auction.pdf>.
- Ausubel, Lawrence M. (1999), "Computer Implemented Methods and Apparatus for Auction," U.S. Patent No. 5,905,975, issued 18 May 1999.
- Ausubel, Lawrence M. (2000), "Computer Implemented Methods and Apparatus for Auction," U.S. Patent No. 6,021,398, issued 1 Feb. 2000.
- Ausubel, Lawrence M. and Paul Milgrom (2001), "System and Method for a Dynamic Auction with Package Bidding," Patent Cooperation Treaty patent application.
- Ausubel, Lawrence M. and Paul Milgrom (2002), "Ascending Auctions with Package Bidding," *Frontiers of Theoretical Economics* **1**(1): Article 1.
<http://www.bepress.com/bejte/frontiers/vol1/iss1/art1>
- Ausubel, Lawrence M. and Paul Milgrom (2004), "The Lovely but Lonely Vickrey Auctions," chapter 1 in this volume.
- Bernheim, B. Douglas and Michael Whinston (1986), "Menu Auctions, Resource Allocation and Economic Influence," *Quarterly Journal of Economics*, **101**: 1-31.
- Bikhchandani, Sushil and Joseph M. Ostroy (2002), "The Package Assignment Model," *Journal of Economic Theory*, **107**: 377-406.
- Brewer, Paul and Charles 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**: 857-86.

- Clarke, E.H. (1971), "Multipart Pricing of Public Goods," *Public Choice*, **11**: 17-33.
- Gale, David and Lloyd S. Shapley (1962), "College Admissions and the Stability of Marriage," *American Mathematical Monthly*, **69**: 9-15.
- Groves, Theodore (1973), "Incentives in Teams," *Econometrica*, **41**: 617-31.
- Hatfield, John W. and Paul Milgrom (2004), "Auctions, Matching and the Law of Aggregate Demand," working paper, Department of Economics, Stanford University.
- Kelso, Alexander S. and Vincent P. Crawford (1982), "Job Matching, Coalition Formation, and Gross Substitutes," *Econometrica*, **50**(6): 1483-1504.
- Milgrom, Paul (2000a), "Putting Auction Theory to Work: The Simultaneous Ascending Auction," *Journal of Political Economy*, **108**(2): 245-272.
- Milgrom, Paul (2000b), "Putting Auction Theory to Work: Ascending Auctions with Package Bidding," unpublished working paper.
- Milgrom, Paul (2004), *Putting Auction Theory to Work*, Cambridge: Cambridge University Press.
- Milgrom, Paul R. and Robert J. Weber (1982), "A Theory of Auctions and Competitive Bidding," *Econometrica*, **50**(5): 1089-1122.
- Parkes, David C. (2001), *Iterative Combinatorial Auctions: Achieving Economic and Computational Efficiency*, Ph.D. thesis, Department of Computer and Information Science, University of Pennsylvania. <http://www.cis.upenn.edu/~dparkes/diss.html>.

Parkes, David C. and Lyle H. Ungar (2000), “Iterative Combinatorial Auctions: Theory and Practice,” *Proceedings of the 17th National Conference on Artificial Intelligence (AAAI-00)*, 74–81.

Roth, Alvin E. and Marilda A. Oliveira Sotomayor (1990), *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Cambridge: Cambridge University Press.

Shapley, Lloyd (1962), “Complements and Substitutes in the Optimal Assignment Problem,” *Naval Research Logistics Quarterly*, **9**: 45-48.

Vickrey, William (1961), “Counterspeculation, Auctions and Competitive Sealed Tenders,” *Journal of Finance*, **16**: 8-37.

Yokoo, Makoto, Yuko Sakurai and Shigeo Matsubara (2004), “The Effect of False-Name Bids in Combinatorial Auctions: New Fraud in Internet Auctions,” *Games and Economic Behavior*, **46**(1): 174-188.

¹ For a more complete account of the pre-auction positioning, see Milgrom (2004), chapter 1.

² Vickrey (1961) originally treated auctions with multiple units of a homogeneous product, while Clarke (1971) and Groves (1973) treated public choice problems. In particular, the Clarke-Groves treatment explicitly includes both the auctions Vickrey studied and auctions of multiple heterogeneous objects. For auction applications, we use “VCG mechanism” and “Vickrey auction” interchangeably.

³ The main ideas underlying the ascending proxy auction and its analysis were introduced in Ausubel (1997, 1999, 2000), Milgrom (2000a,b), Parkes and Ungar (2000), Parkes (2001) and Ausubel and Milgrom (2002).

⁴ Some of our results continue to hold when the bidder retains some discretion to modify its proxy instructions, but we limit attention here to the pure proxy case.

⁵ This formulation rules out “value interdependencies” in which one bidder’s information might affect another bidder’s choices. This assumption would not apply if, for example, the rankings of outcomes

depend on the resolution of common uncertainties about technology or demand. See Milgrom and Weber (1982).

⁶ In the case of public decisions with transfers, as described below, each x_l describes the joint decision and l 's transfer. This allows the bidder to express preferences over a joint decision.

⁷ An allocation x is *stable* if (1) it is feasible ($x \in F$), (2) it is individually rational for each bidder and for the seller, and (3) there exists no S and allocation $y \in F_S$ such that $y \succ_0 x$ and for all $l \in S$, $y_l \succ_l x_l$ or $y_l = x_l$. The notion of blocking used to define the NTU-core replaces (3) with the weaker requirement that there exists no S and allocation $y \in F_S$ such that $y_l \succ_l x_l$ for all $l \in S \cup \{0\}$, ruling out the possibility that $y_l = x_l$ for some $l \in S$.

⁸ See Yokoo, Sakurai and Matsubara (2004) and chapter 7 of this volume for an alternative treatment of skills and “false name bidding.”

⁹ The “semi-sincere” rubric arises because the bidder’s reports of *relative* valuations are truthful, but his reports of *absolute* valuations may be untruthful (i.e., shaded). In the NTU generalization, semi-sincere strategies rank every pair of outcomes accurately except pairs that include the null (non-participation) outcome; semi-sincere strategies may report the null outcome to be higher in the bidder’s ranking than it actually is.

¹⁰ For example, if each bidder raises his bid by the same bid increment whenever he is not a provisional winner, we see that (so long as Bidder 2 remains in the auction), Coalition $\{1, 2\}$ is a provisional winner $\frac{1}{4}$ of the time, Coalition $\{1, 3\}$ is a provisional winner $\frac{1}{4}$ of the time, and Coalition $\{4, 5\}$ is a provisional winner $\frac{1}{2}$ of the time. With starting prices of zero, straightforward bidding would lead Bidder 1 to bid 10, Bidders 2 and 3 to bid 15 each, and Bidders 4 and 5 to bid 10 apiece. At this point, Bidders 4 and 5 drop out of the auction, and so there is nothing further to induce Bidder 1 to raise his bid. But he has already, irrevocably, reached a bid of 10; when his bidder-Pareto-optimal core payment equals 0. If Bidder 1 had instead limited his bidding to 0, he still would have won the West-10 and East-10 licenses.

¹¹ This condition was introduced in Shapley’s (1962) treatment of the assignment problem. It was shown to imply that there is a unique bidder-optimal core point (coinciding with the Vickrey payoff) in Theorem 3 of Ausubel (1997b), which was superceded and improved upon by Theorem 7 of Ausubel and Milgrom

(2002). The predecessor theorem also included the following necessary condition for the uniqueness of a bidder-optimal core point: $w(L) - w(L - S) \geq \sum_{l \in S} (w(L) - w(L - l))$ for all coalitions S ($0 \in S \subset L$). Bikhchandani and Ostroy (2002) subsequently developed the implications of these conditions for dual problems to the package assignment problem.