

The Clock-Proxy Auction: A Practical Combinatorial Auction Design

Lawrence M. Ausubel, Peter Cramton, and Paul Milgrom

Abstract

We propose the clock-proxy auction as a practical means for auctioning many related items. A clock auction phase is followed by a last-and-final proxy round. The approach combines the simple and transparent price discovery of the clock auction with the efficiency of the proxy auction. Linear pricing is maintained as long as possible, but then is abandoned in the proxy round to improve efficiency and enhance seller revenues. The approach has many advantages over the simultaneous ascending auction. In particular, the clock-proxy auction has no exposure problem, eliminates incentives for demand reduction, and prevents most collusive bidding strategies.

1 Introduction

In this chapter we propose a method for auctioning many related items. A typical application is a spectrum sale in which different bidders combine licenses in different ways. Some pairs of licenses may be substitutes and others may be complements. Indeed, a given pair of licenses may be substitutes for one bidder but complements for another, and may change between substitutes and complements for a single bidder as the prices of the other licenses vary. Our proposed method combines two auction formats—the clock auction and the proxy auction—to produce a hybrid with the benefits of both.

The *clock auction* is an iterative auction procedure in which the auctioneer announces prices, one for each of the items being sold. The bidders then indicate the quantities of each item desired at the current prices. Prices for items with excess demand then increase, and the bidders again express quantities at the new prices. This process is repeated until there are no items with excess demand.

The *ascending proxy auction* is a particular package bidding procedure with desirable properties (see Ausubel and Milgrom 2002, Chapter 3). The bidders report values to their respective proxy agents. The proxy agents iteratively submit package bids on behalf of the bidders, selecting the best profit opportunity for a bidder given the bidder's inputted values. The auctioneer then selects the provisionally winning bids that maximize revenues. This process continues until the proxy agents have no new bids to submit.

The clock-proxy auction is a hybrid auction format that begins with a clock phase and ends with a final proxy round. First, bidders directly submit bids in a clock auction, until there is no excess demand for any item. Then bidders have a single opportunity to input proxy values. The proxy round concludes the auction. All bids are kept live throughout the auction. There are no bid withdrawals. The bids of a particular bidder are mutually exclusive. There is an activity rule throughout the clock phase and between the clock phase and the proxy round.

There are three principal motivations behind our clock-proxy auction proposal. First, Porter et al. (2003) precede us in proposing a particular version of a “combinatorial” clock auction for spectrum auctions, and they provide experimental evidence in its support. Second, the recent innovation of the proxy auction provides a combinatorial auction format suitable for related items such as spectrum. Unlike pure clock auctions, whose anonymous linear prices are not generally rich enough to yield efficient outcomes even with straightforward bidding, the proxy auction leads to efficient outcomes and it yields competitive revenues when bidding is straightforward. It also has some desirable individual and group incentive properties. However, the theoretical development of the proxy auction treats only a sealed-bid procedure, omitting opportunities for bidder feedback and price discovery. Third, our own version of a clock auction has been implemented in the field for products such as electricity in recent years with considerable success (see Ausubel and Cramton 2004). This empirical success in the field suggests that the clock phase would be a simple and effective device for providing essential price discovery in advance of a final proxy round. During the clock phase, bidders learn approximate prices for individual items as well as packages (summing the individual prices). This price information helps bidders focus their valuation analysis on packages that are most relevant.

An important benchmark for comparison is the simultaneous ascending auction (see Cramton, Chapter 4; Milgrom 2000, 2004). This auction form performs well when items are substitutes and competition is strong. The clock phase by itself also does well in this simple setting and, in particular, the outcome is similar to that of a simultaneous ascending auction. However, the addition of the proxy auction round should be expected to handle complications, such as complements, collusion, and market power, much better than the simultaneous ascending auction. In environments—including many spectrum auctions—where such complications are present, the clock-proxy auction is likely to outperform the simultaneous ascending auction both on efficiency and revenues.

We begin by motivating and describing the clock phase. Then we examine the proxy phase. Finally we combine the two together in the clock-proxy auction, describing the important role played by both phases, comparing the auction with the simultaneous ascending auction, and discussing implementation issues. Some aspects of the auction technology are further described by Ausubel and Milgrom (2001), Ausubel, Cramton, and Jones (2002), and Milgrom (2004).

2 Clock phase

The simultaneous clock auction is a practical implementation of the fictitious “Walrasian auctioneer.” The auctioneer announces anonymous linear prices. The bidders respond with quantities desired at the specified prices. Then the prices are increased for items in excess demand, while other prices remain unchanged. This process is repeated until there is no excess demand for any item.

The clock phase has several important benefits. First, it is simple for the bidders. At each round, the bidder simply expresses the quantities desired at the current prices. Linear pricing means that it is trivial to evaluate the cost of any package—it is just the inner product of the prices and quantities. Limiting the bidders’ information to a reporting of the excess demand for each item removes much strategizing. Complex bid signaling and collusive strategies are eliminated, as the bidders cannot see individual bids, but only aggregate information. Second, unlike the original Walrasian auction, it is monotonic. This monotonicity contributes to the simplicity of the auction and ensures that it will eventually terminate. Finally, the clock phase

produces highly useable price discovery, because of the item prices (linear pricing). With each bidding round, the bidders get a better understanding of the likely prices for relevant packages. This is essential information in guiding the bidders' decision making. Bidders are able to focus their valuation efforts on the most relevant portion of the price space. As a result, the valuation efforts are more productive. Bidder participation costs fall and efficiency improves.

The weakness of the clock auction is its use of linear pricing at the end of the auction. This means that, to the extent that there is market power, bidders will have an incentive to engage in demand reduction to favorably impact prices. This demand reduction implies that the auction outcome will not be fully efficient (Ausubel and Cramton 2002). When goods are substitutes, the clock auction can restore efficiency by utilizing a "clinching" rule instead of linear pricing (Ausubel 2004, 2006). However, in environments with complementary goods, a clock auction with a separate price quoted for each individual item cannot by itself generally avoid inefficiency. The proxy phase will eliminate this inefficiency.

There are several design choices that will improve the performance of the clock phase. Good choices can avoid the exposure problem, improve price discovery, and handle discrete rounds.

2.1 Avoiding the exposure problem

One important issue in clock auctions is how to treat quantity changes that, if accepted, would make aggregate demand less than supply. For example, for a particular item, demand may equal supply, so the price of the item does not increase, but the increased price of a complementary item may lead the bidder to reduce the quantity it demands. In both clock auctions and the related simultaneous ascending auctions, the usual rule has been to prohibit quantity reductions on items for which the price does not increase, but this creates an exposure problem when some items are complements. Our design allows a bidder to reduce quantity for any item so long as the price has increased on some item the bidder had demanded. This rule eliminates the exposure problem. The bidder is given the flexibility to drop quantity on items for which there is no price increase.

Another case arises when demand is greater than supply for a particular item so the price increases, and one or more bidders attempt to reduce their demands, making demand less than supply. The common approach in this case is to ration the bidders' reductions so that supply equals demand. However, this again creates an exposure problem when some items are complements. Our approach is not to ration the bidders. All reductions are accepted in full.

The reason for the common restrictions on quantity reductions is to avoid undersell (ending the auction at a point where demand is less than supply). However, these restrictions create an exposure problem. Bidders may be forced to purchase quantities that do not make sense given the final price vector. We eliminate these restrictions and avoid the exposure problem. The consequence is the possibility of undersell in the clock phase, but this is of little importance, as the proxy round can resolve any undersell.

We have conducted over twenty high-stake clock auctions using this rule for electricity products, some of which are substitutes and some of which are complements. These are clock-only auctions without a proxy round. However, because the auctions are conducted quarterly, any undersell in the current auction is added to the quantities in the next auction. Our experience has been that undersell typically is slight (only a few percent of the total). The one exception was an auction in which there was a large negative market price shock near the end of the auction, which resulted in undersell of about fifty percent.

With our rule, the clock auction becomes a package auction. For each price vector, the bidder expresses the package of items desired without committing itself to demanding any smaller package.

All bids in the clock phase are kept live in the proxy round. Including these bids has two effects. It potentially increases revenues after the proxy phase by expanding choices in the winner determination problem, and it encourages sincere bidding in the clock phase, because bidders are on the hook for all earlier bids.

2.2 Improving price discovery

In auctions with more than a few items, the sheer number of packages that a bidder might buy makes it impossible for bidders to determine all their values in advance. Bidders adapt to this problem by focusing most of their attention on the packages that are likely to be valuable relative to their forecast prices. A common heuristic to forecast package prices is to estimate the prices of individual items and to take an inner product with quantities to estimate the likely package price. Clock auctions with individual prices assist bidders in this *price discovery* process.

Several recent proposed combinatorial auction procedures, such as the RAD procedure studied by Kwasnica et al. (2005), produce approximate shadow prices on individual items to help guide bidders. The clock auction just does this directly.

Price discovery is undermined to the extent that bidders misrepresent their demands early in the auction. One possibility is that bidders will choose to underbid in the clock phase, hiding as a “snake in the grass” to conceal their true interests from their opponents. To limit this form of insincere bidding, the U.S. Federal Communications Commission (FCC) introduced the Milgrom-Wilson activity rule, and similar activity rules have since become standard in both clock auctions and simultaneous ascending auctions. In its most typical form, a bidder desiring large quantities at the end of the auction must bid for quantities at least as large early in the auction, when prices are lower.

Some clock auctions have performed well in the laboratory without any activity rule (Porter et al. 2003). We suspect that this is because of the limited information that the bidders have about the preferences and plans of the other bidders. This lack of information makes it difficult for participants to know how best to deviate from the straightforward strategy of bidding to maximize profits, ignoring one’s impact on prices. In practice, activity rules appear to be important, because of the more detailed knowledge bidders have about the preferences of others and hence a better sense of the benefits of deviating from straightforward bidding. The first U.S. broadband auction is a good example of an auction where the activity rule played an important role (McAfee and McMillan 1996; Cramton 1997).

The most common activity rule in clock auctions is monotonicity in quantity. As prices rise, quantities cannot increase. Bidders must bid in a way that is consistent with a weakly downward sloping demand curve. This works well when auctioning identical items, but is overly restrictive when there are many different products. If the products are substitutes, it is natural for a bidder to want to shift quantity from one product to another as prices change, effectively arbitraging the price differences between substitute products.

A weaker activity requirement is a monotonicity of aggregate quantity across a group of products. This allows full flexibility in shifting quantity among products in the group. This is the

basis for the FCC’s activity rule. Each license has a number of bidding units associated with it, based on the size of the license. A bidder’s activity in a round is the sum of the bidding units of the licenses on which the bidder is active—either the high bidder in the prior round or placing a valid bid in the current round. This aggregate activity level must exceed or equal a specified percentage (the activity requirement) of the bidder’s current eligibility (typically, 60 percent in the first stage, 80 percent in the second, and 100 percent in the third stage). Otherwise, the bidder’s eligibility in all future rounds is reduced to its activity divided by the activity requirement. Additionally, a bidder has five waivers. A bidder can use a waiver in a round to prevent its eligibility from being reduced in the round.

A weakness of the rule based on monotonicity of aggregate quantities is that it assumes that quantities are readily comparable. For example, in the FCC auctions, the quantity associated with a license is the bandwidth of the license times the population covered (MHz-pop). If prices on a per MHz-pop basis vary widely across licenses, as often is the case, bidders may have an incentive to bid on cheap licenses to satisfy the activity rule. This distortion in bidding compromises price discovery.

We propose an alternative activity rule based on revealed preference that does not require any aggregate quantity measure. The rule is derived from standard consumer theory. Consider any two times, denoted s and t ($s < t$). Let p^s and p^t be the price vectors at these times, let x^s and x^t be the associated demands of some bidder, and let $v(x)$ be that bidder’s value of the package x . A sincere bidder prefers x^s to x^t when prices are p^s :

$$v(x^s) - p^s \cdot x^s \geq v(x^t) - p^s \cdot x^t$$

and prefers x^t to x^s when prices are p^t :

$$v(x^t) - p^t \cdot x^t \geq v(x^s) - p^t \cdot x^s.$$

Adding these two inequalities yields the *revealed preference activity rule*:

$$(RP) \quad (p^t - p^s) \cdot (x^t - x^s) \leq 0.$$

At every time t , the bidder’s demand x^t must satisfy RP for all times $s < t$.

For the case of a single good, RP is equivalent to the condition that as price goes up, quantity cannot increase; that is, bids must be consistent with a weakly downward-sloping demand curve.

Now suppose there are many goods, but all the goods are perfect substitutes in some fixed proportion. For example, the FCC is auctioning 2-MHz licenses and 20-MHz licenses. Ten 2-MHz blocks substitute perfectly for one 20-MHz block. In this simple case, we would want RP to do the same thing it does when the perfect substitutes are auctioned as a single good, and it does so.

First suppose that all prices are consistent with the rate of substitution (e.g., the 20-MHz block is ten times as expensive as the 2-MHz block) and all are increasing by the same percentage. The bidder then only cares about the total quantity in MHz and does not care about which goods are purchased. In this case, RP allows the bidder to substitute arbitrarily across goods. RP is satisfied with equality so long as the bidder maintains the same total MHz in response to the higher prices, and inequality if the bidder reduces total MHz.

Second suppose that the prices are not consistent with the rate of substitution. Say the price on the 2-MHz block increases too fast relative to the 20-MHz block. The bidder then wants to shift all its quantity to the 20-MHz block, and RP allows this: because the 20 MHz is relatively

cheaper, RP gives the bidder more credit for dropping quantity on the 2-MHz blocks than the bidder is debited for the increase in the 20-MHz block. It might seem that the mispricing allows the bidder to expand quantity somewhat, but this is not the case. Because RP is required with respect to all previous bids, the bidder would be constrained by its maximum quantity the last time the 20-MHz block was the best value.

We conclude that RP does just the right thing in the case of perfect substitutes. The activity rule is neither strengthened nor weakened by alternative product definitions.

Now suppose some goods are perfect complements in fixed proportion. For example, in an electricity auction, the bidder wants to maintain a 2-to-1 ratio between baseload product and peakload product. If there are just these two products, then the bidder just cares about the weighted sum of the product prices. As prices increase, the bidder certainly satisfies RP by maintaining the same quantities or by reducing the quantities in the desired ratio; however, the bidder is unable to increase quantities. RP does just the right thing in the case of perfect complements.

If we combine the two cases above so that some goods are perfect substitutes and some are perfect complements, then RP still does the right thing. Bidders will want to shift quantity to the cheapest substitute in building the package of complements. Shifting away from substitute products for which price is increasing too quickly yields a credit that exceeds the debit from shifting toward the relatively cheap product. Hence, this is allowed under RP. Moreover, RP prevents a bidder who always bids on the cheapest among substitutes goods from expanding its quantity of complementary goods as prices rise.

It is useful to compare RP with the current FCC activity rule, which ignores prices and simply looks at aggregate quantity in MHz-pop. “Parking” is the main problem created by the current rule: to maintain flexibility, a bidder has an incentive to bid on underpriced products or low-value products with high quantity, rather than to bid on products that it actually wants to buy. The bidder does this for two reasons: 1) to keep the prices on desired products from increasing too quickly, while maintaining the flexibility to expand demand on products for which competitor demands fall off faster than expected; and 2) to maintain the flexibility to punish a rival by shifting bidding for the rival’s desired markets if the rival bids for the bidder’s desired markets. Thus, parking is motivated by demand reduction and tacit collusion. But the clock implementation mitigates collusion, because bidders see only excess demand; they do not have the information to know when retaliation is needed, where the retaliation should occur, or how to avoid retaliation. And the final proxy round mitigates demand reduction. Hence, we should expect parking to be much less of a problem in the clock implementation.

The greatest damage from parking comes from price distortions that exclude the high-value bidder from winning an item. Under the FCC rule, bidders are most tempted to park on low-price, high-quantity licenses. These prices may get bid up to the point where the efficient winner drops out, because they enable the parking bidder to bid later on other licenses. In contrast, the RP rule does not allow a bidder to increase its quantity for another license unless there is excess demand for the parking license. Thus, parking is only effective when bidding on underpriced goods. But parking on underpriced goods does no harm; it simply serves to increase the price of the underpriced good. Hence, the revealed-preference activity rule has important advantages over the current FCC activity rule.

The revealed-preference activity rule may appear more complex than the FCC rule based on aggregate quantity. However, it still can be displayed in the same simple way on the bidder’s bid

entry screen. As the bid is entered, an activity cell indicates the amount of slack in the tightest RP constraint, and changes to red when the constraint is violated. Moreover, to the extent that the revealed preference activity rule eliminates complex parking strategies, the rule may be simpler for bidders.

2.3 Handling discrete rounds

Although in theory one can imagine implementing an ascending auction in continuous time, this is hardly ever done in practice. Real clock auctions use discrete rounds for two important reasons. First, communication is rarely so reliable that bidders would be willing to be exposed to a continuous clock. A bidder would find it unsatisfactory if the price clock swept past the bidder's willingness to pay because of a brief communication lapse. Discrete rounds are robust to communication problems. Discrete rounds have a bidding window of significant duration, rarely less than ten minutes and sometimes more than one hour. This window gives bidders time to correct any communication problems, to resort to back-up systems, or to contact the auctioneer and have the round extended. Second, a discrete round auction may improve price discovery by giving the bidders an opportunity to reflect between rounds. Bidders need time to incorporate information from prior rounds into a revised bidding strategy. This updating is precisely the source of price discovery and its associated benefits.

An important issue in discrete-round auctions is the size of the bid increments. Larger bid increments enable the auction to conclude in fewer rounds, but the coarse price grid potentially introduces inefficiencies. Large increments also introduce incentives for gaming as a result of the expanded importance of ties. But using small increments, especially in an auction with many clocks, can greatly increase the number of rounds and, hence, the time required to complete the auction. Bidders generally prefer a shorter auction, which reduces participation costs. It also reduces exposure to market price movements during the auction. This is especially relevant in securities and energy auctions for which there are active secondary markets of close substitutes, and for which underlying price movements could easily exceed the bid increments.

Fortunately, it is possible to capture nearly all of the benefits of a continuous auction and still conduct the auction in a limited number of rounds, using the technique of intra-round bids. With intra-round bids, the auctioneer proposes tentative end-of-round prices. Bidders then express their quantity demands in each auction round at all price vectors along the line segment from the start-of-round prices to the proposed end-of-round prices. If, at any time during the round, the prices reach a point at which there is excess supply for some good, then the round ends with those prices. Otherwise, the round ends with the initially proposed end-of-round prices.

Consider an example with two products. The start-of-round prices are 90, 180 and end-of-round prices are 100, 200. The bidder decides to reduce quantity at two price points (40 percent and 60 percent) between the start-of-round and end-of-round prices, as shown below:

Price Point	Product 1		Product 2	
	Price	Quantity	Price	Quantity
0%	90	8	180	4
40%	94	5	188	4
60%	96	5	192	2
100%	100	5	200	2

The auctioneer aggregates all the bids and determines whether any products clear at price points of up to 100 percent. If not, then the process repeats with new end-of-round prices based on excess demand. If one or more products clear, then we find the first product to clear. Suppose the bidder's drop from 8 to 5 at the 40 percent price point causes product 1 to clear, but product 2 has not yet cleared at the 40 percent price point. Then the current round would post at the 40 percent price point. The next round would have start-of-round prices of 94, 188 (the prices at the 40 percent price point) and, perhaps, end-of-round prices of 94, 208. The price of product 1 stops increasing, as there is no longer excess demand.

Following this exact approach means that the clock phase will typically have more rounds than products. This works fine in an environment where there are multiple units of a relatively limited number of products (all of which are assigned the same price). However, this could be an issue in FCC auctions with hundreds of unique licenses requiring independent prices. In that event, the auctioneer may wish to adopt an approach of settling for approximate clearing in the clock phase in order to economize on the number of rounds.

This use of intra-round bids avoids the inefficiency associated with a coarser price grid. It also avoids the gaming behavior that arises from the increased importance of ties with coarser prices. The only thing that is lost is the within-round price discovery. However, within-round price discovery is much less important than the price discovery that occurs between rounds.

The experience from a number of high-stakes clock auctions indicates that intra-round bidding lets the auctioneer conduct auctions with several products in about ten rounds, with little or no loss from the discreteness of rounds (Ausubel and Cramton 2004). These auctions can be completed in a single day. By way of contrast, early spectrum auctions and some electricity auctions without intra-round bids took weeks or even months to conclude. In a few instances, the longer duration was warranted due to the enormous uncertainty and extremely high stakes, but generally speaking, intra-round bids would have reduced the bidding costs without any meaningful loss in price discovery.

2.4 End of the clock phase

The clock phase concludes when there is no excess demand on any item. The result of the clock phase is much more than this final assignment and prices. The result includes all packages and associated prices that were bid throughout the clock phase. Due to complementarities, the clock phase may end with substantial excess supply for many items. If this is the case, the final assignment and prices may not provide a good starting point for the proxy phase. Rather, bids from an earlier round may yield an assignment with higher revenue. (When calculating revenues excess supply should be priced at the reserve price, which presumably represents the seller's opportunity cost of selling the item.)

A sensible approach is to find the revenue maximizing assignment and prices from all the bids in the clock phase. This point is found by backing up the clock to the price point where revenue is at its maximum. The revenue maximizing prices from the clock phase can serve as reasonable lower bounds on prices in the proxy phase. That is, the minimum bid on each package is calculated as the inner product of the revenue maximizing prices and the quantities of items in the package.

In some cases the auctioneer may decide to end the clock phase early—with some excess demand on one or more items. This would be done when the total revenue ceases to increase or

when revenue improvements from successive clock rounds are sufficiently small. With the proxy phase to follow, there is little loss in either revenues or efficiency from stopping, say when revenue improvements are less than one-half percent for two consecutive rounds. At this point price discovery is largely over on all but the smallest items. Giving the auctioneer the discretion to end the clock phase early also enables the auction to follow a more predictable schedule.

3 Proxy phase

Like the clock auction, the proxy auction is based on package bids. However, the incentives are quite different. The main difference is the absence of anonymous linear prices on individual items. Only packages are priced—and the prices may be bidder specific. This weakens price discovery, but the proxy phase is not about price discovery. It is about providing the incentives for efficient assignment. All the price discovery occurs in the clock phase. The second main difference is that the bidders do not bid directly in the proxy phase. Rather, they submit values to the proxy agents, who then bid on their behalf using a specific bidding rule. The proxy agents bid straightforwardly to maximize profits. The proxy phase is a last-and-final opportunity to bid.

The proxy auction works as follows (see Ausubel and Milgrom 2002, Chapter 3). Each bidder reports his values to a proxy agent for all packages that the bidder is interested in. Budget constraints can also be reported. The proxy agent then bids in an ascending package auction on behalf of the real bidder, iteratively submitting the allowable bid that, if accepted, would maximize the real bidder's profit (value minus price), based on the reported values. The auction in theory is conducted with negligibly small bid increments. After each round, provisionally winning bids are determined that maximize seller revenue from compatible bids. All of a bidder's bids are kept live throughout the auction and are treated as mutually exclusive. The auction ends after a round with no new bids (see Hoffman et al., Chapter 17 and Day and Raghavan 2004 for practical methods to implement the proxy phase).

The advantage of this format is that it ends at a core allocation for the reported preferences. Denote the coalition form game (L, w) where L is the set of players ($l = 0$ is the seller and the rest are the bidders) and $w(S)$ is the value of coalition S . Let X denote the set of feasible allocations $(x_i)_{i \in L}$. If S excludes the seller, then $w(S) = 0$; if S includes the seller, then

$$w(S) = \max_{x \in X} \sum_{i \in S} v_i(x_i).$$

The $\text{Core}(L, w)$ is the set of all imputations π (payoffs imputed to the players based on the allocation) that are feasible for the coalition of the whole and cannot be blocked by any coalition S ; that is, for each coalition S , $\sum_{i \in S} \pi_i(x_i) \geq w(S)$.

Theorem (Ausubel and Milgrom 2002, Parkes and Ungar 2000). *The payoff vector π resulting from the proxy auction is a core imputation relative to the reported preferences: $\pi \in \text{Core}(L, w)$.*

Core outcomes exhibit a number of desirable properties, including: 1) efficiency, and 2) competitive revenues for the seller. Thus, the theorem shows that the proxy auction is not subject to the inefficiency of demand reduction: no bidder can ever reduce the price it pays for the package it wins by withholding some of its losing bids for other packages. The theorem also includes the idea that the seller earns competitive revenues: no bidder or coalition of bidders is willing to bid more for the seller's goods. Ausubel and Milgrom (2002, Theorems 2 and 14) establish the core outcome result, whereas Parkes and Ungar (2000, Theorem 1) independently

demonstrate the efficiency of outcomes of an ascending proxy auction without addressing the issue of the core.

A payoff vector in the core is said to be *bidder optimal* if there is no other core allocation that all bidders prefer. If the items are substitutes, then the outcome of the proxy auction coincides with the outcome of the Vickrey auction and with the unique bidder-optimal point in the core. If the goods are not substitutes, then the Vickrey payoff is not generally in the core and the proxy auction yields an outcome with higher seller revenues.

Theorem (Ausubel and Milgrom 2002). *If π is a bidder-optimal point in the $\text{Core}(L,w)$, then there exists a full-information Nash equilibrium of the proxy auction with associated payoff vector π .*

These equilibria may be obtained using strategies of the form: bid your true value minus a nonnegative constant on every package. We emphasize that this conclusion concerns full-information Nash equilibrium: bidders may need to know π to compute their strategies.

Two important advantages of the proxy auction over the Vickrey auction are that the prices and revenues are monotonic (increasing the set of bidders leads to higher prices) and the payoffs are competitive. To illustrate the comparative weaknesses of the Vickrey auction, suppose there are two identical items and two bidders. Bidder 1 values the pair only at \$2.05. Bidder 2 wants a single item only and has a value of \$2. The Vickrey auction awards the pair to bidder 1 for a price of \$2, which is the opportunity cost incurred by not assigning an item to bidder 2. So far, the outcome is unproblematic.

Let us now add a bidder 3 with the same values as bidder 2. In this case, the Vickrey auction awards the items to bidders 2 and 3. Bidder 2's Vickrey price is the opportunity cost of its good to the other participants, which is $\$2.05 - 2.00 = \0.05 . Bidder 3's price is the same. Total revenues fall from \$2.00 to \$0.10. Moreover, the new outcome is not in the core, because the coalition of the seller and bidder 1 could both do better by making a private deal, for example by trading the package at a price of \$1. By way of contrast, adding a bidder in the proxy auction can never reduce seller revenues.

4 The clock-proxy auction

The clock-proxy auction begins with a clock auction for price discovery and concludes with the proxy auction to promote efficiency.

The clock auction is conducted with the revealed-preference activity rule until there is no excess demand on any item. The market-clearing item prices determine the initial minimum bids for all packages for all bidders. Bidders then submit values to proxy agents, who bid to maximize profits, subject to a relaxed revealed-preference activity rule. The bids from the clock phase are kept live as package bids in the proxy phase. All of a bidder's bids, both clock and proxy, are treated as mutually exclusive. Thus, the auctioneer obtains the provisional winning bids after each round of the proxy phase by including all bids—those submitted in the clock phase as well as those submitted in the proxy phase—in the winner determination problem and by selecting at most one provisional winning bid from every bidder. As usual, the proxy phase ends after a round with no new bids.

4.1 Relaxed revealed-preference activity rule

To promote price discovery in the clock phase, the proxy agent's allowable bids must be constrained by the bidder's bids in the clock phase. The constraint we propose is a relaxed version of the revealed preference activity rule.

First, we restate revealed preference in terms of packages and the associated minimum bids for the packages. Consider two times s and t ($s < t$). Suppose the bidder bids for the package S at time s and T at time t . Let $P^s(S)$ and $P^s(T)$ be the package price of S and T at time s ; let $P^t(S)$ and $P^t(T)$ be the package price of S and T at time t ; and let $v(S)$ and $v(T)$ be the value of package S and T . Revealed preference says that the bidder prefers S to T at time s :

$$v(S) - P^s(S) \geq v(T) - P^s(T)$$

and prefers T to S at time t :

$$v(T) - P^t(T) \geq v(S) - P^t(S).$$

Adding these two inequalities yields the revealed preference activity rule for packages:

$$(RP') \quad P^t(S) - P^s(S) \geq P^t(T) - P^s(T).$$

Intuitively, the package price of S must have increased more than the package price of T from time s to time t , for otherwise, at time t , S would be more profitable than T .

Notice that the constraint RP' is automatically satisfied at any two times in the proxy phase, because the proxy agent is required to bid to maximize profits. However, an activity rule based on RP' is too strict when comparing a time s in the clock phase with a time t in the proxy phase. Due to the linear pricing in the clock phase, the bidders have an incentive to reduce demands below their true demands. One purpose of the proxy phase is to let the bidders undo any inefficient demand reduction that would otherwise occur in the clock phase and to defect from any collusive split of the items that would otherwise take place. Hence, it is important to let the bidders expand their demands in the proxy phase. The amount of expansion required depends on the competitiveness of the auction.

We propose a *relaxed revealed-preference activity rule*:

$$(RRP) \quad \alpha[P^t(S) - P^s(S)] \geq P^t(T) - P^s(T).$$

At every time t in the proxy phase, the proxy agent is permitted to bid on the package T only if RRP is satisfied for every package S bid at time s in the clock phase. The proxy agent bids to maximize profits, subject to satisfying RRP relative to all prior bids.

The auctioneer chooses the parameter $\alpha > 1$ based on the competitiveness of the auction. For highly competitive auctions, little demand reduction is likely to occur in the clock phase and α can be set close to 1. On the other hand, if there is little competition (and high concentration), then a higher α is appropriate.

It is possible to state RRP in terms of a restriction on the value function v reported to the proxy, rather than on the bids. Intuitively, a bidder's reported value for a package is constrained by all of its bids in the clock phase. In particular, if the bidder bid on some package S but not T at some time s , then it may not claim at the proxy phase that a bid on T would have been much more profitable, as formalized by the inequality: $v(T) - P^s(T) \leq \alpha(v(S) - P^s(S))$. Under this version of RRP , a bidder is required to state in the proxy phase a value for each package on which the bidder has already bid in the clock phase. The advantage of this approach is that it

allows the proxies to bid accurately according to the bidders' reported values while still imposing consistency across stages.

4.2 Why include the clock phase?

The clock phase provides price discovery that bidders can use to guide their calculations in the complex package auction. At each round, bidders are faced with the simple and familiar problem of expressing demands at specified prices. Moreover, because there is no exposure problem, bidders can bid for synergistic gains without fear. Prices then adjust in response to excess demand. As the bidding continues, bidders get a better understanding of what they may win and where their best opportunities lie.

The case for the clock phase relies on the idea that it is costly for bidders to determine their preferences. The clock phase, by providing tentative price information, helps focus a bidder's decision problem. Rather than consider all possibilities from the outset, the bidder can instead focus on cases that are important given the tentative price and assignment information. Although the idea that bidders can make information processing decisions in auctions is valid even in auctions for a single good (Compte and Jehiel 2000), its importance is magnified when there are many goods for sale, because the bidder's decision problem is then much more complicated. Rather than simply decide whether to buy at a given price, the bidder must decide which goods to buy and how many of each. The number of possibilities grows exponentially with the number of goods. Price discovery can play an extremely valuable role in guiding the bidder through the valuation process.

Price discovery in the clock phase makes bidding in the proxy phase vastly simpler. Without the clock phase, bidders would be forced either to determine values for all possible packages or to make uninformed guesses about which packages were likely to be most attractive. Our experience with dozens of bidders suggests that the second outcome is much more likely; determining the values of exponentially many packages becomes quickly impractical with even a modest number of items for sale. Using the clock phase to make informed guesses about prices, bidders can focus their decision making on the most relevant packages. The bidders see that they do not need to consider the vast majority of options, because the options are excluded by the prices established in the clock phase. The bidders also get a sense of what packages are most promising, and how their demands fit in the aggregate with those of the other bidders.

In competitive auctions where the items are substitutes and competition is strong, we expect the clock phase to do most of the work in establishing prices and assignments—the proxy phase would play a limited role. When competition is weak, demand reduction may lead the clock phase to end prematurely, but this problem is corrected at the proxy stage, which eliminates incentives for demand reduction. If the clock auction gives the bidders a good idea of likely package prices, then expressing a simple approximate valuation to the proxy is made easier. For example, with global economies of scope, a bidder might report to his proxy bidder a value for each item, a fixed cost of operation, and a limit on the number of items acquired. This is just an example, but it serves to highlight that simple valuation functions might serve well once the range of likely package prices is limited.

4.3 Why include the proxy phase?

The main advantage of the proxy phase is that it pushes the outcome toward the core, that is, toward an efficient allocation with competitive payoffs for the bidders and competitive revenues for the seller.

In the proxy phase, there are no incentives for demand reduction. A large bidder can bid for large quantities without the fear that doing so will adversely impact the price the bidder pays.

The proxy phase also mitigates collusion. Any collusive split of the items established in the clock phase can be undone in the proxy phase. The relaxed activity rule means that the bidders can expand demands in the proxy phase. The allocation is still up for grabs in the proxy phase.

The clock-proxy auction has some similarities with the Anglo-Dutch design initially proposed for (but not ultimately used in) the United Kingdom's third-generation mobile wireless auction (Klemperer 2002). Both formats have an ascending auction followed by a sealed-bid last-and-final round. However, the motivation for the last-and-final round is quite different. In the Anglo-Dutch design, the last round has pay-as-bid pricing intended to introduce inefficiency, so as to motivate inefficient bidders to participate in the auction (and perhaps increase auction revenues). In the clock-proxy auction, the last round is more similar to Vickrey pricing and is intended to promote efficiency, rather than prevent it. The relaxed activity rule in the proxy round, however, does encourage the undoing of any tacit collusion in the clock phase, and in this sense is similar to the last-and-final round of the Anglo-Dutch design.

The proxy phase will play a more important role to the extent that competition is limited and complementarities are strong and varied across bidders. Then it is more likely that the clock phase will end prematurely. However, in competitive auctions, the proxy phase may not be needed.

A potential problem with a clock-only auction under our proposed rules arises from a bidder's ability to reduce quantity on products even when the price of a product does not go up. This may appear to create a "free withdrawal" and a potential source of gaming. For example, a bidder might bid up a competitor on a competitor's preferred license to the point where the competitor drops out. Then the strategic bidder reduces quantity on this product. Alternatively, the bidder might bid up the competitor and then drop quantity before the competitor drops out.

Two features mitigate this potential problem. First, the revealed-preference activity rule makes it risky for a bidder to overbid on items that the bidder does not want. Unlike the activity rule based on aggregate quantity, the bidder dropping quantity on a product for which the price has not increased is not given any credit in the RP inequality and hence has no ability to expand demand on another product. Second, the preferred approach would run the winner-determination-problem at the end among *all* prior bids. Hence, the strategic bidder may find that it is obligated to purchase items that it does not want. (Of course, if goods are mostly substitutes, then one simply could prevent quantity reductions for goods that have cleared.)

4.4 Two examples

We illustrate our answers to "Why include the clock phase?" and "Why include the proxy phase?" with two examples.

In our first example, there are two items and two bidders. Bidder 1 wants just a single item and values it at v_1 . Bidder 2 wants up to two items and values each at v_2 (valuing the package of

two items at $2v_2$). The private values v_1 and v_2 are drawn independently from the uniform distribution on $[0,1]$. Each bidder i knows the realization of v_i but only the distribution of v_j ($j \neq i$). In the clock auction, this is a classic example of demand reduction. For simplicity, assume that the clock price ascends continuously. Bidder 1's weakly dominant strategy is to bid a quantity of 1 at all prices up to v_1 and then to drop to a quantity of 0. Bidder 2 has a choice whether to bid initially for a quantity of two, or to bid for only one unit and cause the price clock to stop at zero. A straightforward calculation shows that bidding for only one unit and obtaining a zero price maximizes bidder 2's expected payoff, establishing that this is the unique equilibrium (Ausubel and Cramton 2002, p. 4).

Thus, conducting only a clock phase is disastrous for the seller; revenues equal zero and the outcome of each bidder winning one unit is inefficient whenever $v_2 > v_1$. However, suppose that the clock phase is followed by a proxy round, using a parameter $\alpha \geq 2$ in the relaxed revealed-preference activity rule. Because the substitutes condition is satisfied in this example, the bidders' dominant strategies in the proxy round are each to bid their true values. Thus, the clock-proxy auction yields the bidder-optimal core outcome, and the seller earns revenues of $\min\{v_1, v_2\}$. Nothing of consequence occurs in the clock phase, and the proxy phase yields the desirable outcome by itself.

In our second example, there are m items and n bidders ($n > m$). Each bidder i values item k at v_{ik} . But bidder i has value for only a single item, and so for example if bidder i received both items k and l , his value would be only $\max\{v_{ik}, v_{il}\}$. The values v_{ik} are random variables with support $[0,1]$. Each bidder i knows the realization of v_{ik} ($k = 1, \dots, m$), but only the distribution of v_{jk} ($j \neq i$) ($k = 1, \dots, m$). In the clock auction, because bidders have demand for only a single item, each bidder's dominant strategy is to bid a quantity of one on an item k such that $v_{ik} - p_k = \max_{l=1, \dots, m} \{v_{il} - p_l\}$ and to bid a quantity of zero on all other items. Therefore, the clock phase concludes at the Vickrey outcome, which is also the predicted outcome of the proxy phase (because the substitutes condition is satisfied). Thus, the clock-proxy auction again yields the bidder-optimal core outcome. This time the clock phase yields the desirable outcome by itself, and nothing further occurs in the proxy phase.

If the bidders find it costly to determine their values, the clock phase may find the outcome without the need for bidders to calculate all their values. For example, suppose $m = 2$ and $n = 3$ and the bidders' estimated value pairs are $(2,4)$, $(3,8)$ and $(7,2)$, but each bidder knows each of its values only to within ± 1 , without further costly investment. In the clock phase, bidder 1 will be the first to face the need to invest in learning its exact values. If he does so, the auction will end at prices of 2 and 4 without the second and third bidder ever needing to make that investment. Price discovery at the clock phase saves bidders 2 and 3 from the need to determine their full values for the proxy stage.

4.5 Comparison with the simultaneous ascending auction

The simultaneous ascending auction as implemented by the FCC is an important benchmark of comparison, given its common use in auctioning many related items (see Cramton, Chapter 4). The clock auction is a variant of the simultaneous ascending auction in which the auctioneer

specifies prices and the bidders name quantities. There are several advantages to the clock implementation.

The clock auction is a simpler process than the simultaneous ascending auction. Bidders are provided the minimal information needed for price discovery—the prices and the excess demand. Bidders are not distracted by other information that is either extraneous or useful as a means to facilitate collusion.

The clock auction also can take better advantage of substitutes, for example, using a single clock for items that are near perfect substitutes. In spectrum auctions, there is a tendency for the spectrum authority to make specific band plans to facilitate the simultaneous ascending auction. For example, anticipating demands for a large, medium, and small license, the authority may specify a band plan with three blocks—30 MHz, 20 MHz, and 10 MHz. Ideally, these decisions would be left to the bidders themselves. In a clock auction, the bidders could bid the number of 2-MHz blocks desired at the clock price. Then the auction would determine the band plan, rather than the auction authority. This approach is more efficient and would likely be more competitive, because all bidders are competing for all the bandwidth in the clock auction. With the preset band plan, some bidders may be uninterested in particular blocks, such as those that are too large for their needs.

Clock auctions are faster than a simultaneous ascending auction. Simultaneous ascending auctions are especially slow near the end, when there is little excess demand. For example, when there are six bidders bidding on five similar licenses, then it typically takes five rounds to obtain a one bid-increment increase on all items. In contrast, in a clock auction, an increment increase takes just a single round. Moreover, intra-round bids allow larger increments, without introducing inefficiencies, because bidders still can express demands along the line segment from the start-of-round prices to the end-of-round prices.

The clock auction limits collusion relative to the simultaneous ascending auction. Signaling how to split up the items is greatly limited. Collusive strategies based on retaliation are not possible, because bidder-specific quantity information is not given. Further, the simultaneous ascending auction can have a tendency to end early when an obvious split is reached, but this cannot happen in the clock auction, because the bidders lack information about the split. Also there are fewer rounds to coordinate a split.

The clock auction, as described here, eliminates the exposure problem. As long as at least one price increases, a bidder can reduce quantity on his other items. The bid is binding only as a full package. Hence, the bidder can safely bid for synergistic gains.

The clock-proxy auction shares all these advantages of the clock auction, and in addition promotes core outcomes. The proxy phase further mitigates collusion and eliminates demand reduction. The cost of the proxy phase is added implementation complexity. Also the absence of linear pricing reduces the transparency of the auction. It is less obvious to a bidder why he lost. Nonetheless, the auctioneer at the conclusion of the auction can disclose sufficient information for the bidders to determine the outcome without revealing any supramarginal values.

4.6 Combinatorial exchange

Like other package auctions, the clock-proxy auction is designed for settings with a single seller. With multiple sellers and no item prices, there is an additional problem to solve: how to divide the auction revenues. For example, if separate sellers own items A and B, and if all the

bidders want to buy items A and B together, with no interest in these separate and separately owned items, the auction itself can provide no information about how to allocate the revenue from the winning bid among the sellers. The revenue-sharing rule has to be determined separately, and there is no simple and completely satisfactory solution to this problem.

The clock-proxy auction can be extended to handle exchanges with one passive seller and many active buyers and sellers. A natural application is the auctioning of encumbered spectrum (Cramton, Kwerel, and Williams 1998; Kwerel and Williams 2002). The spectrum authority would be the passive seller, selling overlay licenses. Incumbents are (potentially) the active sellers, selling their existing rights. In this setting, one can adapt the clock-proxy auction very simply. An incumbent seller's bid would reflect an offer to sell a package. Formally, its bid would specify the goods it offers as negative quantities in the clock phase and would specify negative quantities and prices in the proxy stage. In principle, one could even allow bids in which an incumbent offers to exchange its good for another good plus or minus some compensating payment, where the package is expressed by a vector of positive and negative numbers.

Alternative designs differ in how they divide auction revenues and in what bids sellers are allowed to make. For example, one possibility is to fix the items to be sold at the proxy stage as those that were not acquired by their original owners at the clock stage. Final revenues would then be distributed to sellers in proportion to the prices from the clock stage. Another possibility is to allow the sellers to bid in every stage of the auction, essentially negotiating what is sold and how revenues are to be split through their bidding behavior. A third possibility is to allow sellers to set reserve prices and to use those to divide revenues among the sellers.

These alternative designs split revenues differently, so they create different incentives for incumbents to report exaggerated values. The result will be differences in the likelihood of a successful sale. So far, theory provides little guidance on which choice is best, beyond indicating that the problem can sometimes be a hard one. If there are many sellers whose goods are sufficiently good substitutes, then the problem may not be too severe. This strongly suggests that the most important issue for the FCC in making the package exchange a success is careful attention to the incumbents' rights, to make their goods as substitutable as possible.

4.7 Implementation issues

We briefly discuss four of the most important implementation issues.

Confidentiality of values

One practical issue with the proxy phase is confidentiality of values. Bidders may be hesitant to bid true values in the proxy phase, fearing that the auctioneer would somehow manipulate the prices with a "seller shill" to push prices all the way to the bidders' reported values. Steps need to be taken to assure that this cannot happen. A highly transparent auction process helps to assure that the auction rules are followed. Auction software can be tested and certified that it is consistent with the auction rules. At the end of the auction, the auctioneer can report all the bids. The bidders can then confirm that the outcome was consistent with the rules. In addition, there is no reason that the auctioneer needs to be given access to the high values. Only the computer need know.

A further step to protect the privacy of high values is to allow a multi-round implementation of the proxy phase. The critical feature of the proxy phase is that the relative values are locked. If

bidders do not want to reveal their final values, that can be handled. In a multi-round version of the proxy phase, bidders must freeze the relative values of the packages they name but can periodically authorize a fixed dollar increase in all of their bids. With this approach, the auction becomes an ascending, pay-as-bid package auction.

Price increments in the clock phase

When auctioning many items, one must take care in defining the price adjustment process. This is especially true when some goods are complements. Intuitively, undersell in the clock phase is minimized by having each product clear at roughly the same time. Otherwise price increases on complementary products can cause quantity drops on products that have already cleared. Thus, the goal should be to come up with a price adjustment process that reflects relative values as well as excess demand. Moreover, the price adjustment process effectively is resolving the threshold problem by specifying who should contribute what as the clock ticks higher. To the extent that prices adjust with relative values the resolution of the threshold problem will be more successful.

One simple approach is to build the relative value information into the initial starting prices. Then use a percentage increase, based on the extent of excess demand. For example, the percentage increment could vary linearly with the excess demand, subject to a lower and upper limit.

Expression of proxy values

Even with the benefit of the price discovery in the clock phase, expressing a valuation function in the proxy phase may be difficult. When many items are being sold, the bidder will need a tool to facilitate translating preferences into proxy values. The best tool will depend on the circumstances.

At a minimum, the tool will allow an additive valuation function. The bidder submits a demand curve for each item. The value of a package is then found by integrating the demand curve (adding the marginal values) up to the quantity of the item in the package, and then adding over all items. This additive model ignores all value interdependencies across items; it assumes that the demand for one item is independent of the demand for other items. Although globally (across a wide range of quantities) this might be a bad assumption, locally (across a narrow range of quantities) this might be a reasonable approximation. Hence, provided the clock phase has taken us close to the equilibrium, so the proxy phase is only doing some fine-tuning of the clock outcome, then such a simplistic tool may perform reasonably well. And of course it performs very well when bidders actually have additive values.

A simple extension of the additive model allows the bidder to express perfect substitutes and complements within the additive structure. For example, items A and B may be designated perfect complements in the ratio 1 to 3 (one unit of A is needed for three units of B). Then the bidder expresses a demand curve for A and B (with the one-to-three ratio always maintained). Items C and D may be designated perfect substitutes in the ratio 2 to 1 (two Cs equal one D). Then the bidder expresses a demand curve for C or D (with all quantity converted to C-equivalent). This extension effectively allows the bidder to redefine the items in such a way to make the additive model fit. For example, in a spectrum auction, a bidder for paired spectrum will want to express a demand for paired spectrum. This can be done by designating the upper and lower channels as perfect complements, but then the blocks of paired spectrum as perfect

substitutes. A bidder for unpaired spectrum would designate all channels as perfect substitutes, and then express a single demand curve for unpaired spectrum.

Demand curves typically are expressed as step functions, although in some contexts piecewise linear demand curves are allowed. Bidders should be able to specify whether quantity can be rationed. For example if a bidder drops quantity from 20 to 10 at a price of \$5, does this mean the bidder is just as happy getting 14 units as 10 units or 20 units when the price is \$5 per unit, or does the bidder only want exactly 10 units at a price of \$5, and exactly 20 units at a price of \$4.99? Is there a minimum quantity that the bidder must win for the item to have value?

Beyond this, the tool should allow for the inclusion of bidder constraints. Budget constraints are the most common: do not bid more than X. Other constraints may be on quantities: only value A if you win B. This constraint arises in spectrum auctions when a bidder has secondary regions that have value only if the primary regions are won.

The bidders' business plans are a useful guide to determine how best to structure the valuation tool in a particular application. Business plans are an expression of value to investors. Although the details of the business plans are not available to the auctioneer, one can construct a useful valuation tool from understanding the basic structure of these business plans.

Calculating prices in the proxy phase

The proxy phase is a sealed-bid auction. At issue is how best to calculate the final assignment and prices. The final assignment is easy. This is just the value maximizing assignment given the reported values. The harder part is determining the prices for each winning package. The clock phase helps by setting a lower bound on the price of each package. Given these starting prices, one approach would be to run directly the proxy auction with negligible bid increments. With many items and bidders this would require voluminous calculations.

Fortunately, one can accelerate the process of calculating prices using various methods (see Hoffman et al., Chapter 17; Day and Raghavan 2004; Zhong et al. 2003). First, as David Parkes suggested, package prices for all bidders can start at "safe prices," defined as the maximum bid on the package by any losing bidder. Second, prices can increase in discrete jumps to the point where a bidder starts or stops bidding on a particular package. Although these methods have not yet been fully developed, calculating the prices in the proxy phase likely can be done with many items and bidders in an expedient manner.

The precise process for calculating the prices is especially important when some items are complements, because then there will be a set of bidder-optimal points in the core, and the price process will determine which of these points is selected.

5 Conclusion

We propose the clock-proxy auction for auctioning many related items—a simultaneous clock auction followed by a last-and-final proxy round. The basic idea is to use anonymous linear prices as long as possible to maximize price discovery, simplicity, and transparency. The clock phase also greatly facilitates the bidders' valuation analysis for the proxy round, because the analysis can be confined to the relevant part of the price space identified in the clock phase. Finally, unlike the simultaneous ascending auction, the clock auction does not suffer from the exposure problem.

For highly competitive auctions of items that are mostly substitutes, the clock auction without the proxy round will perform well. Indeed a clock auction without a proxy round may be the best approach in this setting, as it offers the greatest simplicity and transparency, while being highly efficient.

With limited competition or items with a complex and varied structure of complements, adding the proxy phase can improve the auction outcome. In particular, a core outcome is achieved. Seller revenues are competitive and the allocation is efficient. The demand reduction incentive present in the clock phase is eliminated. Most importantly, adding the proxy round does no harm: in the simplest settings where the clock auction alone performs well, adding the proxy round should not distort the outcome. The proxy round simply expands the settings in which the auction performs well.

Acknowledgments

This research was inspired by the Federal Communications Commission's efforts to develop a practical combinatorial auction for its spectrum auctions. We are especially grateful to Evan Kwerel for his insights and encouragement.

References

- Ausubel, Lawrence M. (2004), "An Efficient Ascending-Bid Auction for Multiple Objects," *American Economic Review*, 94:5, 1452-1475.
- Ausubel, Lawrence M. (2006), "An Efficient Dynamic Auction for Heterogeneous Commodities," *American Economic Review*, forthcoming.
- Ausubel, Lawrence M. and Peter Cramton (2002), "Demand Reduction and Inefficiency in Multi-Unit Auctions," University of Maryland Working Paper 9607, revised July 2002.
- Ausubel, Lawrence M. and Peter Cramton (2004), "Auctioning Many Divisible Goods," *Journal of the European Economic Association*, 2, 480-493, April-May.
- Ausubel, Lawrence M., Peter Cramton, and Wynne P. Jones (2002). "System and Method for an Auction of Multiple Types of Items." International Patent Application No. PCT/ US02/16937.
- Ausubel, Lawrence M. and Paul Milgrom (2001), "System and Method for a Dynamic Auction with Package Bidding," International Patent Application No. PCT/US01/43838.
- Ausubel, Lawrence M. and Paul Milgrom (2002), "Ascending Auctions with Package Bidding," *Frontiers of Theoretical Economics*, 1, 1-45, <http://www.bepress.com/bejte/frontiers/voll/iss1/art1>.
- Compte, Olivier and Philippe Jehiel (2000), "On the Virtues of the Ascending Price Auction." Working paper, CERAS-ENPC.
- Cramton, Peter (1997), "The FCC Spectrum Auctions: An Early Assessment," *Journal of Economics and Management Strategy*, 6:3, 431-495.
- Cramton, Peter, Evan Kwerel, and John Williams (1998), "Efficient Relocation of Spectrum Incumbents," *Journal of Law and Economics*, 41, 647-675.
- Day, Robert W. and S. Raghavan (2004), "Generation and Selection of Core Outcomes in Sealed-Bid Combinatorial Auctions," Working Paper, University of Maryland.
- Klemperer, Paul (2002), "What Really Matters in Auction Design," *Journal of Economic Perspectives*, 16:1, 169-189.
- Kwasnica, Anthony M., John O. Ledyard, Dave Porter, and Christine De Martini (2005), "A New and Improved Design for Multi-Object Iterative Auctions," *Management Science*, forthcoming.
- Kwerel, Evan R. and John R. Williams (2002), "A Proposal for the Rapid Transition to Market Allocation of Spectrum," Working Paper, Office of Plans and Policy, FCC.
- McAfee, R. Preston and John McMillan (1996), "Analyzing the Airwaves Auction," *Journal of Economic Perspectives*, 10, 159-176.
- Milgrom, Paul (2000), "Putting Auctions Theory to Work: The Simultaneous Ascending Auction," *Journal of Political Economy*, 108(2): 245-272.
- Milgrom, Paul (2004), *Putting Auction Theory to Work*, Cambridge: Cambridge University Press.
- 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.
- Porter, David, Stephen Rassenti, Anil Roopnarine, and Vernon Smith (2003), "Combinatorial Auction Design," *Proceedings of the National Academy of Sciences*, 100, 11153-11157.
- Zhong, Jie, Gangshu Cai, and Peter R. Wurman (2003), "Computing Price Trajectories in Combinatorial Auctions with Proxy Bidding," Working Paper, North Carolina State University.