"Auctions: Theory"
For the New Palgrave, 2nd Edition

Lawrence M. Ausubel
University of Maryland

Auctions occupy a deservedly prominent place within microeconomics and game theory, for at least three reasons:

- The auction is, in its own right, an important device for trade. Auctions have long been a common way of selling diverse items such as works of art and government securities. In recent years, their importance in consumer markets has increased through the ascendancy of eBay and other Internet auctions. At the same time, the use of auctions for transactions between businesses has expanded greatly, most notably in the telecommunications, energy and environmental sectors, and for procurement purposes generally.
- Auctions have become the clearest success story in the application of game theory to economics. In most applications of game theory, the modeler has considerable (perhaps excessive) freedom to formulate the rules of the game, and the results obtained will often be highly sensitive to the chosen formulation. By way of contrast, an auction will typically have a well-defined set of rules, yielding clearer theoretical predictions.
- There has been an increasing wealth of auction data available for empirical analysis in recent years. In conjunction with the available theory, this has led to a growing body of empirical work on auctions. Moreover, auctions are very well suited for laboratory experiments and they have been a very fruitful area for experimental economics.

This entry is limited in its scope to *auction theory*. Other related entries include "Auctions: Empirics," which reviews the growing body of empirical work relating to auctions, and "Mechanism Design," which reviews a theoretical body of work that originated in the analysis of problems related to optimal auctions but extends well beyond.

1. Introduction

Auction theory is often said to have originated in the seminal 1961 article by William Vickrey. While Vickrey's insights were initially unrecognized and it would be many years before his work was followed up by other researchers, it eventually led to a formidable body of research by pioneers including Wilson, Clarke, Groves, Milgrom, Weber, Myerson, Maskin and Riley. The first wave of theoretical research into auctions concluded in the mid-1980's, by which time there was a widespread sense that it had become a relatively complete body of work with very little remaining to be discovered. See McAfee and McMillan (1987) for an excellent review of the first wave of auction theory.

However, the perception that auction theory was complete began to change following two pivotal events in the 1990's: the Salomon Brothers scandal in the US Government securities market in 1991; and the advent of the Federal Communications Commission (FCC) spectrum auctions in 1994. In the aftermath of the former, the Department of the Treasury sought input from academia concerning the US Treasury auctions. In the preparation for the latter, the FCC encouraged the active involvement of auction theorists in the design of the new auctions.

Each of these two episodes undoubtedly benefitted from the participation of academics. In particular, the FCC introduced an innovative dynamic auction format — the simultaneous ascending auction — whose empirical performance appears far superior to previous static sealed-bid auctions. The Treasury's experimentation with, and eventual adoption of, uniform-price auctions in place of pay-as-bid auctions, also appears to have resulted from economists' input.

At the same time, these two pivotal events underscored some extremely serious limitations in auction theory as it existed in the early- to mid-1990's. It became apparent then that the theory that had been developed was almost exclusively one of single-item auctions, and that relatively little was established concerning multi-item auctions. As the flip side of the same coin, these episodes made obvious that many of the empirically-important examples of auctions involve a multiplicity of items. As a result, a second wave of theoretical research into auctions, focusing especially on multi-item auctions, emerged in the middle of the 1990's and continues today.

This entry begins by reviewing the theory of single-item auctions, largely completed during the first period of research. It then continues by reviewing the theory of multi-unit auctions, still a work in progress.

The scope and detail of the present article is necessarily quite limited. For deeper and more comprehensive treatments of auctions, three recent and notable books, by Krishna (2002), Milgrom (2004) and Cramton, Shoham and Steinberg (2006), are especially recommended to readers. Earlier survey articles by McAfee and McMillan (1987) and Wilson (1992) also provide excellent treatments of the literature on single-item auctions. A compendium by Klemperer (2000) brings together many of the best articles in auction theory.

2. Sealed-Bid Auctions for Single Items

Much of the analysis within traditional auction theory has concerned sealed-bid auctions (i.e., static games) for single items. Bidders submit their sealed bids in advance of a deadline, without knowledge of any of their opponents' bids. After the deadline, the auctioneer unseals the bids and determines a winner. The two most commonly studied sealed-bid formats are:

FIRST-PRICE AUCTION. The highest bidder wins the item, and pays the amount of his bid.

SECOND-PRICE AUCTION. The highest bidder wins the item, and pays the amount bid by the second-highest bidder.

Note that the above auction formats (and, indeed, all of the auctions described in this entry) have been described for a regular auction in which the auctioneer offers items for sale and the bidders are buyers. Each can easily be restated for a "reverse auction" (i.e., procurement auction) in which the auctioneer solicits the purchase of items and the bidders are sellers. For example, in a second-price reverse auction, the lowest bidder is chosen to provide the item and is paid the amount bid by the second-lowest bidder.

A. The Private Values Model

A seller wishes to allocate a single unit of a good or service among n bidders ($i = 1, \dots, n$). The bidders bid simultaneously and independently as in a noncooperative static game. Bidder i's payoff from receiving the item in return for the payment y is given by $v_i - y$ (whereas bidder i's payoff from not winning the item is normalized to zero). Each bidder i's valuation, v_i, for the item is private information. Bidder i knows v_i at the time he submits his bid. Meanwhile, the opposing bidders $j \neq i$ view v_i as a random variable whose realization is unknown, but which is drawn according to the known joint distribution function $\hat{F}(v_1, \dots, v_i, \dots, v_n)$.

This model is referred to as the *private values* model, on account that each bidder's valuation depends only on his own — and not the other bidders' — information. (By contrast, in a *pure common values* model, $v_i = v_j$, for all $i, j = 1, \dots, n$; and in an *interdependent values* model, bidder i's valuation is allowed to be a function of $v_{-i} = \{v_j\}_{j \neq i}$, as well as of v_i.) With private values, some especially simple and elegant results hold, particularly for the second-price auction.

Two additional assumptions are frequently made. First, we generally assume that bidders are *risk neutral* in evaluating their payoffs under uncertainty. That is, each bidder seeks merely to maximize the mathematical expectation of his payoff. Second, we often assume *independence* of the private information. That is, the joint distribution function, $\hat{F}(v_1, \ldots, v_n)$, is given by the product of separate distribution functions, $F_i(\cdot)$, for each of the v_i. However, both the risk neutrality and independence assumptions are unnecessary for solving the second-price auction, which we analyze first.

B. Solution of the Second-Price Auction

Sincere bidding (i.e., the truthful bidding of one's own valuation) is a Nash equilibrium of the sealed-bid second-price auction, under private values. That is, if each bidder *i* submits the bid $b_i = v_i$, then there is no incentive for any bidder to unilaterally deviate. Moreover, sincere bidding is a weakly dominant strategy for each bidder; and sincere bidding by all bidders is the unique outcome of elimination of weakly dominated strategies. These facts make the sincere bidding equilibrium an especially compelling outcome of the second-price auction.

Let $\hat{b}_{-i} = \max_{j \neq i} \{b_j\}$, the highest among the opponents' bids. The dominant strategy property is easily established by comparing bidder *i*'s payoff from the sincere bid of $b_i = v_i$ with his payoff from instead bidding $b_i{}' < v_i$ ("shading" his bid). If $\hat{b}_{-i}$ is less than $b_i{}'$ or greater than v_i, then bid-shading has no effect on bidder *i*'s payoff; in the former case, bidder *i* wins either way, and in the latter case, bidder *i* loses either way. However, in the event that $\hat{b}_{-i}$ is between $b_i{}'$ and v_i, the bid-shading makes a difference: if bidder *i* bids v_i, he wins the auction and thereby achieves a positive payoff of $v_i - \hat{b}_{-i} > 0$; whereas, if bidder *i* bids $b_i{}'$, he loses the auction and receives zero payoff. Thus, $b_i = v_i$ weakly dominates any bid $b_i{}' < v_i$. A similar comparison finds that $b_i = v_i$ weakly dominates any bid $b_i{}' > v_i$. Sincere bidding is optimal, regardless of the bidding strategies of opposing bidders.

Note that the above argument in no way uses the risk neutrality or independence assumptions; nor does it require any form of symmetry. Another way of viewing this is that sincere bidding is an *ex post equilibrium* of the second-price auction, in the sense that the strategy would remain optimal even if the bidder were to learn his opponents' bids before he was required to submit his own bid. Indeed, one of the strengths of the result that sincere bidding is a Nash equilibrium in weakly dominant strategies is that it basically relies only upon the private values assumption, and is otherwise extremely robust to the specification of the model.

C. Incentive Compatibility in any Sealed-Bid Auction Format

Consider any equilibrium of *any* sealed-bid auction format, in the private values model. Given that bidder *i*'s valuation is private information, observe that there is nothing to force bidder *i* to bid according to his true valuation v_i instead of some other valuation w_i. As a result, the equilibrium must have a structure that gives bidder *i* the incentive to bid according to his true valuation. This requirement is known as *incentive compatibility*.

In the following derivation, we assume that the support of each bidder *i*'s valuation is the interval $[\underline{v}_i, \overline{v}_i]$. We will make both the risk-neutrality and independence assumptions. Let $\Pi_i(v_i)$ denote bidder *i*'s expected payoff, let $P_i(v_i)$ denote bidder *i*'s probability of winning the item, and let $Q_i(v_i)$ denote bidder *i*'s expected payment in this equilibrium, when his valuation is v_i. The reader should note that $Q_i(v_i)$ refers here to bidder *i*'s unconditional expected payment — *not* to his expected payment conditional on winning. Given the risk-neutrality assumption, $\Pi_i(v_i)$ is given by:

(1) $$\Pi_i(v_i) = P_i(v_i)\, v_i - Q_i(v_i) \;.$$

Next, we pursue the observation that there is nothing forcing bidder *i* to bid according to his true valuation v_i rather than according to another valuation w_i. Define $\pi_i(w_i, v_i)$ to be bidder *i*'s expected payoff from employing the bidding strategy of a bidder with valuation w_i when his true valuation is v_i. Observe that:

(2) $$\pi_i(w_i, v_i) = P_i(w_i)\, v_i - Q_i(w_i) \;,$$

since bidder *i*'s probability of winning and expected payment depend exclusively on his bid, not on his true valuation. Bidder *i* will voluntarily choose to bid according to his true valuation only if his expected payoff is greater than from bidding according to another valuation w_i, that is, if:

(3) $$\Pi_i(v_i) \geq \pi_i(w_i, v_i), \text{ for all } v_i, w_i \in [\underline{v}_i, \overline{v}_i] \text{ and all } i = 1, \ldots, n \;.$$

Inequality (3), referred to as the *incentive-compatibility constraint*, has very strong implications.

Next, note that $\Pi_i(v_i) = \pi_i(v_i, v_i) = \max_{w_i \in [\underline{v}_i, \overline{v}_i]} \pi_i(w_i, v_i)$. It is straightforward to see that $\Pi_i(\cdot)$ is monotonically nondecreasing and continuous. Consequently, it is differentiable almost everywhere and equals the integral of its derivative. Applying the Envelope Theorem at any v_i where $\Pi_i(\cdot)$ is differentiable yields:

$$\text{(4)} \qquad \frac{d\Pi_i(v_i)}{dv_i} = \left.\frac{\partial \pi_i(w_i, v_i)}{\partial v_i}\right|_{w_i=v_i} = P_i(w_i)\big|_{w_i=v_i} = P_i(v_i) \ .$$

Integrating Eq. (4), we have:

$$\text{(5)} \qquad \Pi_i(v_i) = \Pi_i(\underline{v}_i) + \int_{\underline{v}}^{v_i} P_i(x)dx \, , \text{ for all } v_i \in [\underline{v}_i, \overline{v}_i] \text{ and all } i = 1, \ldots, n \, .$$

D. Solution of the First-Price Auction

The sealed-bid first-price auction requires two symmetry assumptions in order to yield a fairly simple solution. First, we assume *symmetric bidders*, in the sense that the joint distribution function $\hat{F}(v_1, \ldots, v_i, \ldots, v_n)$ governing the bidders' valuations is a symmetric function of its arguments. This assumption and the associated notation are simplest to state if independence is assumed. In this case, we write $F_i(\cdot)$ for the distribution function of each v_i; symmetry is the assumption that $F_i = F$, for all $i = 1, \ldots, n$, or in other words, the assumption that the various v_i are identically distributed, as well as independent, random variables. However, a similar derivation with only slightly more cumbersome notation is possible if the bidders are symmetric but the v_i are affiliated random variables. We write $[\underline{v}, \overline{v}]$ for the support of $F(\cdot)$. In addition, we assume that $F(\cdot)$ is a continuous function, so that there are no mass points in the common probability distribution of the bidders' valuations.

Second, we restrict attention to *symmetric, monotonically increasing equilibria* in pure strategies. The assumed symmetry of bidders opens the possibility for existence of a symmetric equilibrium. (Meanwhile, asymmetric equilibria are also possible in symmetric games, but Maskin and Riley, 2003, establish that, under slightly stronger assumptions, the construction here gives the unique equilibrium of the auction.) Any pure-strategy equilibrium can be characterized by the bid functions $\{B_i(\cdot)\}_{i=1}^n$, which give bidder *i*'s bid $B_i(v_i)$ when his valuation is v_i. Our assumption is that $B_i = B$, for all $i = 1, \ldots, n$, where $B(\cdot)$ is a strictly increasing function.

Observe that in any symmetric equilibrium, bidder *i* wins against bidder *j* if and only if $B(v_j) < B(v_i)$ and, given strict monotonicity, if and only if $v_j < v_i$. (We can ignore the event $v_j = v_i$; this is a zero-probability event, since we have assumed the distribution of valuations has no mass points.) Consequently, bidder *i* wins the item if and only $v_j < v_i$ for all $j \neq i$. Since the

$\{v_j\}_{j\neq i}$ are i.i.d. random variables, bidder *i* has probability $F(v_i)^{n-1}$ of winning the auction when his valuation is v_i. We write: $P_i(v_i) = F(v_i)^{n-1}$, for all $v_i \in [\underline{v}, \overline{v}]$ and all $i = 1, \ldots, n$.

Moreover, in a first-price auction, the bidder's payoff equals $v_i - B(v_i)$ if he wins the auction and zero if he loses. Consequently his expected payoff equals:

(6) $$\Pi_i(v_i) = P_i(v_i)\left[v_i - B(v_i)\right] = F(v_i)^{n-1}\left[v_i - B(v_i)\right].$$

Observe from Eq. (6) that, if $v_i = \underline{v}$, bidder *i*'s probability of winning equals zero and, hence, $\Pi_i(\underline{v}) = 0$. Substituting this fact and $P_i(v_i) = F(v_i)^{n-1}$ into Eq. (5) yields:

(7) $$\Pi_i(v_i) = \int_{\underline{v}}^{v_i} F(x)^{n-1} dx, \text{ for all } v_i \in [\underline{v}, \overline{v}] \text{ and all } i = 1, \ldots, n.$$

Combining Eq. (6) with Eq. (7), and solving for $B(\cdot)$, yields the equilibrium bid function:

(8) $$B(v_i) = v_i - \frac{\Pi_i(v_i)}{F(v_i)^{n-1}} = v_i - \frac{\int_{\underline{v}}^{v_i} F(x)^{n-1} dx}{F(v_i)^{n-1}}.$$

The posited strict monotonicity is verified by differentiating Eq. (8) with respect to v_i, which shows that $B'(v_i) > 0$. Thus, Eq. (8) provides us with the unique symmetric equilibrium in pure strategies of the sealed-bid first-price auction. This result holds for arbitrary continuous distribution functions $F(\cdot)$ with support on an interval $[\underline{v}, \overline{v}]$.

3. Revenue Equivalence, Efficient Auctions and Optimal Auctions

Standard practice in auction theory is to evaluate auction formats according to either of two criteria: efficiency and revenue optimization. With the quasi-linear utilities generally assumed in auction theory, efficiency means putting the items in the hands of those who value them the most. Revenue maximization means maximizing the seller's expected revenues or, in a procurement auction, minimizing the buyer's expected procurement costs. In auctions of government assets such as spectrum licenses, the explicit objective is often efficiency. In auctions by private parties, the explicit objective is often revenue optimization.

A. Efficient Auctions

The above solutions to the second-price and first-price auctions both yield full efficiency. In the symmetric increasing equilibrium of the first-price auction, the highest bid corresponds to the highest valuation, and so the item is assigned efficiently for every realization of the random variables. In the dominant strategy equilibrium of the second-price auction, the identical conclusion holds. Thus, in a symmetric private values model, an objective of efficiency looks kindly upon both auction formats — but does not prefer one over the other.

B. Revenue Equivalence

One of the classic and most far-reaching results in auction theory is Revenue Equivalence, which provides a set of assumptions under which the seller's and buyers' expected payoffs are guaranteed to be the same under different auction formats.

Revenue Equivalence (Vickrey, 1961; Myerson, 1981; Riley and Samuelson, 1981) may be stated as follows. Assume that the random variables representing the bidders' valuations are independent, and assume that bidders are risk neutral. Consider any two auction formats satisfying both of the following properties: (1) The two auction formats assign the item(s) to the same bidder(s), for every realization of random variables; and (2) the two auction formats give the same expected payoff to the lowest valuation type, $\underline{v}_i$, of each bidder *i*. Then each bidder earns the same expected payoff under each of the two auction formats and, consequently, the seller earns the same expected revenues under each of the two auction formats.

For an auction of a single item, the result follows directly from Eq. (5) above. Recall that this equation holds for any equilibrium of *any* sealed-bid auction format. If for every realization of the random variables, the two auction formats assign the item to the same bidder, then each bidder's probability, $P_i(\cdot)$, of winning is the same under the two auction formats. If in addition, $\Pi_i(\underline{v}_i)$ is the same under the two auction formats, then Eq. (5) implies that the entire function $\Pi_i(\cdot)$ is the same under the two auction formats. Since this holds for every bidder *i*, and since the expected gains from trade are the same under the two auction formats, it follows from an accounting identity that the seller's expected revenues are also the same under the two auction formats.

One of the most important applications of Revenue Equivalence is that the above solutions to the second-price and first-price auctions give the seller the same expected revenues (and also give each buyer the same expected payoffs). Revenue equivalence is applicable because, as argued above, the item is assigned efficiently for every realization of the random variables in each of these auction formats. Moreover, when $v_i = \underline{v}$, the expected payoff of bidder *i* equals zero in each of these auction formats. To understand this result, observe that (all other things equal) a bidder in a first-price auction will bid lower than in a second-price auction, since the payment rule is less generous. Expected revenues will be greater in the first-price or the second-price auction depending on whether the highest of a collection of smaller bids or the second-highest of a collection of larger bids is greater in expectation. The Revenue Equivalence theorem establishes that, in the symmetric private values model, the two effects exactly offset one another.

C. Optimal Auctions

Another classic result of auction theory is the determination of the auction format that optimizes revenues. This result, known in the literature as the *optimal auction*, is due to Harris and Raviv (1981), Myerson (1981), and Riley and Samuelson (1981). Any possible auction format is considered — the item may be assigned to the bidder who submitted the highest bid (as in the second-price or first-price auction), but it may alternatively be allocated to another bidder, randomized in its allocation, or withheld from sale entirely, depending on the collection of bids submitted. At the outset, this might be viewed as a very complicated problem, since it requires selecting simultaneously the probability of winning and a payment that optimizes revenues. However, using analysis similar to the treatment of incentive compatibility, above, it can be shown that the expected payment is determined up to a constant by the probability of winning. Consequently, the problem simplifies to determining the probability of each bidder winning (for every realization of the random variables) that optimizes revenues.

For symmetric bidders, each of whose distributions satisfy a regularity condition, a particularly simple characterization of the optimal auction can be obtained. Let $F(\cdot)$ be the distribution function of the valuation v_i of each bidder *i*, let $f(\cdot)$ be the associated density function and suppose that $v_i - \frac{1-F(v_i)}{f(v_i)}$ is strictly increasing in v_i for all $v_i \in [\underline{v}, \overline{v}]$. Then the optimal auction assigns the item to the bidder *i* with the highest v_i, if and only if the highest v_i exceeds the reserve valuation *r*, where *r* is defined by $r - \frac{1-F(r)}{f(r)} = v_0$ and where v_0 is the seller's valuation for the item.

In other words, with symmetric bidders, both the second-price and the first-price auctions become optimal auctions, once a reserve price of *r* is inserted.

D. Full Rent Extraction

The optimal auctions problem can be reconsidered without the independence assumption. However, Crémer and McLean (1985) demonstrate that, if the bidders' private information is correlated, then there exists a mechanism that enables the seller to extract *all* of the gains from trade. The mechanism includes a procedure for allocating the item efficiently. Superimposed on this, the mechanism provides rewards to bidders if their reports of private information "agree" with each other and penalties to bidders if their reports "disagree" with each other. The amounts of the rewards and penalties — both potentially quite large in magnitude — are set so as to make the bidders indifferent between participating and not participating in the mechanism. As such, the mechanism enables the seller to extract the entire surplus, including the informational rents that the bidders are able to obtain under the independence assumption. This is referred to as *full rent extraction*.

Crémer and McLean's result may be viewed as fundamentally negative, in that it suggests that the optimal auctions analysis may be of limited relevance. Real-world auction mechanisms appear to be broadly consistent with the predictions of the optimal auctions theory under the independence assumption, but they look nothing like the full rent-extracting mechanisms possible with correlated private information. Given that there are good reasons to believe that bidders' private information is correlated with one another, it would appear that the optimal auctions analysis does not provide us with great insight into real-world auctions. Some subsequent research has attempted to weaken the extreme conclusion of full rent extraction by positing that bidders have limited liability or by introducing opportunities for auctioneer collusion or cheating, but in many respects these devices appear to be ineffectual patches for an elegant theory (optimal auctions) that suffers from only limited empirical relevance.

4. Dynamic Auctions for Single Items

The next two formats considered for auctioning single items are dynamic auctions: participants bid sequentially over time and, potentially, learn something about their opponents' bids during the course of the auction. In the first dynamic auction, the price *ascends*; and in the second dynamic auction, the price *descends*:

English Auction. Bidders dynamically submit successively higher bids for the item. The final bidder wins the item, and pays the amount of his final bid.

Dutch Auction. The auctioneer starts at a high price and announces successively lower prices, until some bidder expresses his willingness to purchase the item by bidding. The first bidder to bid wins the item, and pays the current price at the time he bids.

Note that, as in Section 2, each of these auction formats has been described for a regular auction in which the auctioneer offers items for sale, but can easily be restated for a "reverse auction." For example, in an English reverse auction, the bids would descend rather than ascend; while in a Dutch reverse auction, the auctioneer would offer to buy at successively higher prices.

A. Solution of the Dutch Auction

An insight due to Vickrey (1961) is that the Dutch auction is strategically equivalent to the sealed-bid first-price auction. To see the equivalence, consider the real meaning of a strategy b_i by bidder i in the Dutch auction: "If no other bidder bids for the item at any price higher than b_i, then I am willing to step in and purchase it at b_i." Just as in the sealed-bid first-price auction, the bidder i who selects the highest strategy b_i in the Dutch auction wins the item and pays the amount b_i. Furthermore, although the Dutch auction is explicitly dynamic, there is nothing that can happen that would lead any bidder to want to change his strategy while the auction is still running. If strategy b_i was a best response for bidder i evaluated at the starting price p_0, then b_i remains a best response evaluated at any price $p < p_0$, assuming that no other bidder has already bid at a price between p_0 and p. Meanwhile, if another bidder has already bid, then there is nothing that bidder i can do; the Dutch auction is over. Hence, any equilibrium of the sealed-bid first-price auction is also an equilibrium of the Dutch auction, and vice versa.

B. Solution of the English Auction

By way of contrast, some meaningful learning and/or strategic interaction is possible during an English auction, so the outcome is potentially different from the outcome of the sealed-bid second-price auction.

We model the English auction as a "clock auction": the auctioneer starts at a low price and announces successively higher prices. At every price, each bidder is asked to indicate his willingness to purchase the item. The price continues to rise so long as two or more bidders indicate interest. The auction concludes at the first price such that fewer than two bidders indicate interest, and the item is awarded at the final price. This clock-auction description is used instead of a game where bidders successively announce higher prices, since it yields simpler arguments and clean results.

With pure private values, the reasonable equilibrium of the English auction corresponds to the dominant-strategy equilibrium of the sealed-bid second-price auction. A bidder's strategy designates the price at which he will drop out of the auction (assuming that at least one opponent still remains); in equilibrium, the bidder sets his drop-out price equal to his true valuation. However, matters become more complicated in the case of interdependent valuations, where each bidder's valuation depends not only on his own information, v_i, but also on the opposing bidders' information, v_{-i}. We turn to this case next.

C. The Winner's Curse and Revenues under Interdependent Values

One of the most celebrated phenomena in auctions is the "Winner's Curse". Whenever a bidder's valuation depends positively on other bidders' information, winning an item in an auction may confer "bad news" in the sense that it indicates that other bidders possessed adverse information about the item's value. The potential for falling victim to the Winner's Curse may induce restrained bidding, curtailing the seller's revenues. In turn, some auction formats may produce higher revenues than others, to the extent that they mitigate the Winner's Curse and thereby make it safe for bidders to bid more aggressively.

The basic intuition, which is often referred to as the "linkage principle" and is due to Milgrom and Weber (1982), is that the Winner's Curse is mitigated to the extent that the winner's payment depends on the opposing bidders' information. Thus, under appropriate assumptions, the second-price auction will yield higher expected revenues than the first-price auction: the price paid by the winner of a second-price auction depends on the information possessed by the highest losing bidder, while the price paid by the winner of a first-price auction depends exclusively on his own information. Moreover, the English auction will yield higher expected revenues than the second-price auction: the price paid by the winner of an English auction may depend on the information possessed by *all* of the losing bidders (who are observed as they drop out), while the price paid by the winner of a (sealed-bid) second-price auction depends only on the information of the *highest* losing bidder.

These conclusions require an assumption known as "affiliation," which intuitively means something very close to "nonnegative correlation." More precisely, let $v = (v_1, \dots, v_n)$ and $v' = (v'_1, \dots, v'_n)$ be possible realizations of the n bidders' random variables, and let $f(\cdot, \dots, \cdot)$ denote the joint density function. Let $v \vee v'$ denote the component-wise maximum of v and v', and let $v \wedge v'$ denote the component-wise minimum. The random variables v and v' are said to be *affiliated* if:

(9) $$f(v \vee v')\, f(v \wedge v') \geq f(v)\, f(v'), \text{ for all } v, v' \in [\underline{v}, \dots, \overline{v}]^n .$$

Affiliation provides that two high realizations or two low realizations of the random variables are at least as likely as one high and one low realization, etc., meaning something close to nonnegative correlation. Independence is included (as a boundary case) in the definition: for independent random variables, the affiliation inequality (9) is satisfied with equality. To obtain strict revenue rankings, the affiliation inequality must hold strictly.

These conclusions also rely on several symmetry assumptions. Bidders are symmetric, the equilibria considered are symmetric, and each bidder's valuation depends on all of its opponents' information in a symmetric way. Each bidder's valuation increases (weakly) in its own and its opponents' information, and attention is restricted to equilibria in monotonically increasing strategies. As before, each bidder is risk neutral in evaluating its payoff under uncertainty.

These conclusions also rely on a monotonicity assumption: each bidder's valuation increases (weakly) in its own and in the opposing bidders' information. In addition, as before, each bidder is risk-neutral in evaluating its payoff under uncertainty. Furthermore, the two symmetry assumptions of Section 2D are made: bidders are symmetric, in the sense that the joint distribution governing the bidders' information is a symmetric function of its arguments; and attention is restricted to symmetric, monotonically increasing equilibria in pure strategies.

Under these assumptions, the sealed-bid first-price and second-price auctions and the English auction possess symmetric, monotonic equilibria. However, while these equilibria are all efficient, Milgrom and Weber (1982) establish that they may be ranked by revenues: the English auction yields expected revenues greater than or equal to those of the sealed-bid second-price auction, which in turn yields expected revenues greater than or equal to those of the sealed-bid first-price auction. Their theorem provides one of the most powerful results of auction theory, justifying the conventional wisdom that dynamic auctions yield higher revenues than sealed-bid auctions.

5. Auctions of Homogeneous Goods

A. Sealed-Bid, Multi-Unit Auction Formats

The defining characteristic of a homogeneous good is that each of the *M* individual items is identical (or a close substitute), so that bids can be expressed in terms of quantities without indicating the identity of the particular good that is desired. Treating goods as homogeneous has the effect of dramatically simplifying the description of the bids that are submitted and the overall auction procedure. This simplification is especially appropriate in treating subject matter such as financial securities or energy products. Any two $10,000 US government bonds with the same interest rate and the same maturity *are* identical, just as any two megawatts of electricity provided at the same location on the electrical grid at the same time *are* identical.

There are three principal sealed-bid, multi-unit auction formats for M homogeneous goods. In each of these, a bid comprises an inverse demand function, i.e., a (weakly) decreasing function $p_i(q)$, for $q \in [0, M]$, representing the price offered by bidder i for a first, second, etc. unit of the good. (Note that this notation may be used to treat situations where the good is perfectly divisible, as well as situations where the good is offered in discrete quantities.) The bidders submit bids; the auctioneer then aggregates the bids and determines a clearing price. Each bidder wins the quantity demanded at the clearing price, but his payment varies according to the particular auction format:

PAY-AS-BID AUCTION.[1] Each bidder wins the quantity demanded at the clearing price, and pays the amount that he bid for each unit won.

UNIFORM-PRICE AUCTION.[2] Each bidder wins the quantity demanded at the clearing price, and pays the clearing price for each unit won.

MULTI-UNIT VICKREY AUCTION. Each bidder wins the quantity demanded at the clearing price, and pays the opportunity cost (relative to the bids submitted) for each unit won.

Sealed-bid, multi-unit auction formats are best known in the financial sector for their longtime and widespread use in the sale of government securities. For example, a survey of OECD countries in 1992 found that Australia, Canada, Denmark, France, Germany, Italy, Japan, New Zealand, the United Kingdom and, of course, the United States then used sealed-bid auctions for selling at least some of their debt. The pay-as-bid auction was the traditional format used for government securities, and it was used for all US Treasury bills until a decade ago. The uniform-price auction was first proposed seriously as a replacement for the pay-as-bid auction by Milton Friedman in testimony at a 1959 Congressional hearing. Wilson (1979) gave the first theoretical analysis of a uniform-price auction. In 1993, the US began an "experiment" of using

[1] Pay-as-bid auctions are also known as "discriminatory auctions" or "multiple-price auctions."

[2] Uniform-price auctions are often referred to in the financial press as "Dutch auctions," generating some confusion with respect to the standard usage of the auction theory literature. They are also known as "nondiscriminatory auctions," "competitive auctions," or "single-price auctions."

the uniform-price auction for two- and five-year government notes and, beginning in 1998, the US switched entirely to the uniform-price auction for all issues. Meanwhile, the multi-unit Vickrey auction was introduced and first analyzed in Vickrey's 1961 paper.

The pay-as-bid auction can be correctly viewed as a multi-unit generalization of the first-price auction. However, it is quite difficult to calculate Nash equilibria of the pay-as-bid auction, unless efficient equilibria exist. *Three* symmetry assumptions together guarantee the existence of efficient equilibria. First, bidders are assumed to be symmetric, in the sense that the joint distribution governing the bidders' information is symmetric with respect to the bidders. Second, bidders regard every unit of the good as symmetric: that is, each bidder *i* has a constant marginal valuation for every quantity $q_i \in [0, \lambda_i]$, up to a capacity of λ_i, and a marginal valuation of zero thereafter. Third, the bidders are symmetric in their capacities: that is, $\lambda_i = \lambda$, for all bidders *i*. With these assumptions, the pay-as-bid auction has a solution very similar to that of the first-price auction for a single item. However, without these assumptions, it inherits an undesirable property from the single-item auction: absent symmetry, all Nash equilibria of the pay-as-bid auction will generally be inefficient (Ausubel and Cramton, 2002, Theorems 3 and 4).

The uniform-price auction bears a superficial resemblance to the second-price auction of a single item, in that a high winning bid gains the benefit of a lower marginal bid. However, any similarity is indeed only superficial as, except under very restrictive assumptions, all equilibria of the uniform-price auction are inefficient. The argument is simplest in the same model of constant marginal valuations as in the previous paragraph. If the capacities of all bidders are equal (i.e., if $\lambda_i = \lambda$ for all *i*) and if the supply is an integer multiple of λ, then there exists an efficient Bayesian-Nash equilibrium of the uniform-price auction. (For example, if there are *M* identical units available and if every bidder has a unit demand, then sincere bidding is a Nash equilibrium in dominant strategies.) However, if if the bidders' capacities are unequal or if the supply is not an integer multiple of λ, then all equilibria of the uniform-price auction are inefficient (Ausubel and Cramton, 2002, Theorems 2 and 5).

The intuition for inefficiency in the uniform-price auction can be found by taking a close look at optimal bidding strategies. Sincere bidding is weakly dominant for a *first* unit: if a bidder's first bid determines the clearing price, then the bidder wins zero units. However, the bidder's *second* bid may determine the price he pays for his first unit, providing an incentive to shade his bid. The extent of *demand reduction*, as this bid shading is known, increases in the number of units, since the number of inframarginal units whose price may be affected increases.

Further, note that the allocation rule in the auction has the effect of equating the amounts of the bidders' marginal bids. Since a large bidder will likely have shaded his marginal bid more than a small bidder, the large bidder's marginal value is probably greater than a small bidder's. Consequently, the bidders' marginal values will be unequal, contrary to efficiency.

Meanwhile, the Vickrey auction is the correct multi-unit generalization of the second-price auction. As in the pay-as-bid and uniform-price auctions, bidders simultaneously submit inverse demand functions and each bidder wins the quantity demanded at the clearing price. However, rather than paying the bid price or the clearing price for each unit won, a winning bidder pays the *opportunity cost*. If a bidder wins K units, he pays the K^{th} highest rejected bid of his opponents for his first unit, the $(K-1)^{\text{st}}$ highest rejected bid of his opponents for his second unit, … , and the highest rejected bid of his opponents for his K^{th} unit. The dominant strategy property of the sealed-bid second-price auction generalizes because a bidder's payment is determined solely by his opponents' bids. Consequently, given pure private values and nonincreasing marginal values, sincere bidding is an efficient equilibrium in weakly dominant strategies.

B. Efficiency and Revenue Comparisons

Under pure private values, the dominant strategy equilibrium of the Vickrey auction attains full efficiency. It can be shown that neither the pay-as-bid nor the uniform-price auction generally attains efficiency; moreover, the efficiency ranking of these two formats is inherently ambiguous. Continuing the argument of the previous subsection, it is sufficient to examine environments in which bidders have constant marginal valuations. If $F_i = F$ and $\lambda_i = \lambda$ for all bidders *i*, but the supply is *not* an integer multiple of λ, then the pay-as-bid auction has an efficient equilibrium while all equilibria of the uniform-price auction are inefficient. Conversely, if $\lambda_i = \lambda$ for all bidders *i* and if the supply is an integer multiple of λ, but $F_i \neq F_j$ for two bidders *i* and *j*, then the uniform-price auction has an efficient equilibrium while all equilibria of the pay-as-bid auction are generally inefficient (Ausubel and Cramton, 2002).

On revenues, the policy literature has generally assumed that the uniform-price auction outperforms the pay-as-bid auction; however, the argument of the previous paragraph can be extended to reverse the assumed ranking. Maskin and Riley (1989) extend Myerson's (1981) characterization of the optimal auction to multiple homogeneous goods: with symmetric bidders and constant marginal valuations, their characterization requires allocating items efficiently. Thus, as in the previous paragraph, if $F_i = F$ and $\lambda_i = \lambda$ for all bidders *i*, but the supply is *not* an integer multiple of λ, then the efficient equilibrium of the pay-as-bid auction outranks all equilibria of the uniform-price auction on *revenues* (as well as efficiency).

C. Uniform-Price Clock Auctions

The "clock auction" — a practical design for dynamic auctions of one or more types of goods, with its origins in the "Walrasian auctioneer" from the classical economics literature — has seen increasing use as a trading institution since 2001. A fictitious auctioneer is often presented as a device or thought experiment for understanding convergence to a general equilibrium. The Walrasian auctioneer announces a price vector, *p*; bidders report the quantity vectors that they wish to transact at these prices; and the auctioneer increases or decreases each component of price according as excess demand is positive or negative (*Walrasian tâtonnement*). This iterative process continues until a price vector is reached at which excess demand is zero, and trades occur only at the final price vector. In real-world applications, instead of a fictitious auctioneer serving as a metaphor for a market-clearing process, the process is taken literally; a real auctioneer announces prices and accepts bids of quantities. Applications, to date, have largely been in the electricity, natural gas, and environmental sectors.

The basic clock auction differs from the standard Sotheby's or eBay auction in that bidders do not propose prices. Rather, the auctioneer announces prices, and bidders' responses are limited to the reporting of quantities desired at the announced prices, until clearing is attained. As such, it is closest to the auction-theorist's depiction of the English auction for a single item (or the traditional Dutch auction) — but generalized, so that instead of bidders merely giving binary responses of whether they are "in" or "out" as prices ascend, they indicate their quantities desired.

Observe that the uniform-price clock auction is correctly viewed as a dynamic version of the sealed-bid uniform-price auction reviewed in the previous two subsections. The important difference is that, in the dynamic auction, bidders will typically receive repeated feedback as to the aggregate demand at the various prices.

As such, the clock auction may inherit the advantages that dynamic auctions have over sealed-bid auctions, First, under conditions that can be made precise, the insight from single-item auctions that feedback about other bidders' valuations would ameliorate the Winner's Curse and lead to more aggressive bidding carries over to the multi-unit environment. Second, clock auctions, better than sealed-bid auctions, allow bidders to maintain the privacy of their valuations for the items being sold. Bidders never need to submit any indications of interest at any prices beyond the auction's clearing price. Third, when there are two or more types of items, auctioning them simultaneously enables bidders to submit bids based on the substitution possibilities or complementarities among the items at various price vectors. At the same time, the iterative nature of the auction economizes on the amount of information submitted: demands do not need to be submitted for all price vectors, but only for price vectors reached along the convergence path to equilibrium.

Unfortunately, the uniform-price clock auction also inherits the demand reduction and inefficiency of the sealed-bid uniform-price auction. Indeed, as a theoretical proposition, the problem of bidders optimally reducing their quantities bid well below their true demands can become substantially worse in the dynamic version of the auction. The *reductio ad absurdum* is provided by Ausubel and Schwartz (1999), who analyze a two-bidder clock auction game of complete information in which the bidders alternate in their moves. For a wide set of environments, the unique subgame perfect equilibrium has the qualitative description that, at the first move, the first player reduces his quantity to approximately half of the supply and, at the second move, the second player reduces his quantity to clear the market. Thus, the outcome is inefficient and the revenues barely exceed the starting price.

As a practical matter, demand reduction may not undermine the outcome of a uniform-price clock auction where there is substantial competition for every item being sold. However, if one or more of the bidders has considerable market power, it may become important to use an auction format which avoids creating incentives for demand reduction.

D. Efficient Clock Auctions

Ausubel (2004, 2006) proposes an alternative clock auction design, which utilizes the same general structure as the uniform-price clock auction, but adopts a different payment rule that eliminates the incentives for demand reduction. In essence, the design provides a dynamic version of the (multi-unit) Vickrey auction, and thereby inherits its incentives for truth-telling.

The Ausubel auction is easiest described for a homogeneous good. After each set of bidder reports, the auctioneer determines whether any bidder has "clinched" any of the units offered (i.e., whether any bidder is mathematically guaranteed to win one or more units). For example, in an auction with a supply of 5 units, and three bidders demanding 3, 2 and 2 units, respectively, the first bidder has clinched 1 unit, as his opponents' total demand of 4 is less than the supply of

5. Rather than awarding units only at a final uniform price, the auction awards units at the current price whenever they are newly clinched.

If this alternative clock auction is represented as a static auction, it collapses to the Vickrey auction in the same sense that an English auction collapses to the sealed-bid second-price auction. Consequently, it can be proven that sincere bidding is an equilibrium and, in a suitable discrete specification of the game under incomplete information, sincere bidding is the unique outcome of iterated elimination of weakly-dominated strategies. Thus, unlike the uniform-price clock auction, there is no incentive for demand reduction.

6. Auctions of Heterogeneous Goods

In many significant applications, the multiple items offered within an auction are each unique, so it is not adequate for bidders merely to indicate the quantities that they desire. For example, an FCC spectrum auction might include a New York license, a Washington license and a Los Angeles license. Moreover, there might be synergies in owning various combinations: for example, a New York and a Washington license together might be worth more together than the sum of their values separately. Such environments pose particular challenges for auction theory.

A. Simultaneous Ascending Auctions

The simultaneous ascending auction, proposed in comments to the FCC by Paul Milgrom, Robert Wilson and Preston McAfee, has been used in auctions on six continents allocating more than $100 billion worth of spectrum licenses. Some of the best known applications of the simultaneous ascending auction include: the Nationwide Narrowband Auction (July 1994), the first use of the simultaneous ascending auction; the PCS A/B Auction (December 1994 – March 1995), the first large-scale auction of mobile telephone licenses, which raised $7 billion; the United Kingdom UMTS Auction (March – April 2000), which raised 22.5 billion British pounds; and the German UMTS Auction (July – August 2000), which raised 50 billion euro.

In the simultaneous ascending auction, multiple items are put up for sale at the same time and the auction concludes simultaneously for all of the items. As such, it is a modern version of the "silent auction" that is frequently used in fundraisers by charitable institutions. Bidders submit bids in a sequence of rounds. Each bid comprises a single item and an associated price, which must exceed the standing high bid by at least a minimum bid increment. After each round, the new standing high bids for each item are determined. The auction concludes after a round passes in which no new bids are submitted, and the standing high bids are then deemed to be winning bids. Payments equal the amounts of the winning bids.

The critical innovation in the simultaneous ascending auction is the inclusion of *activity rules* into the auction design. Activity rules are bidding constraints that limit a bidder's bidding activity in the current round based on his past bidding activity (i.e., his standing high bids and new bids). Without activity rules, bidders would tend to wait as "snakes in the grass" until nearly the end of the auction before placing their serious bids, thwarting any price discovery (the main reason for conducting a dynamic auction in the first place). Conversely, activity rules have the effect of forcing bidders to place meaningful bids in early rounds of the auction and thereby to reveal information to their opponents.

B. Walrasian Equilibria as Outcomes of Simultaneous Ascending Auctions

A Walrasian equilibrium — consisting of prices for the various items and an allocation of the items to the bidders such that each item with a non-zero price is assigned to exactly one bidder and such that each bidder prefers his assigned allocation to any alternative bundle at the given prices — is a plausible outcome for the simultaneous ascending auction. Assuming that a Walrasian equilibrium was reached, no bidder would have any incentive to attempt to upset the allocation, even if he believed he could obtain additional items without further increasing their prices. Thus, it becomes interesting to identify the conditions needed for existence of Walrasian equilibria with discrete items.

Kelso and Crawford (1982) show that the substitutes condition is sufficient for the existence of Walrasian equilibrium. "Substitutes" literally refers to the price-theoretic condition that if the price of one item is increased while the price of every other item is held fixed, then the demand for every other item weakly increases. Moreover, the substitutes condition is "almost necessary" for existence. Suppose that the set of possible bidder preferences includes all valuation functions satisfying the substitutes condition, but also includes at least one valuation function violating the substitutes condition. Then if there are at least two bidders, there exists a profile of valuation functions such that no Walrasian equilibrium exists (Gul and Stacchetti, 1999, and Milgrom, 2000).

The reader should avoid losing sight of the fact that, just because a Walrasian equilibrium exists for a discrete environment, it does not necessarily follow that the simultaneous ascending auction will terminate at a Walrasian equilibrium. The strongest statement that can be made is that, if bidders bid "straightforwardly" (i.e., if they demand naïvely the bundle of items that maximizes their utility, while ignoring strategic considerations), then a Walrasian equilibrium will be reached. However, observe that even with homogeneous goods, consumers with weakly diminishing marginal valuations satisfy the substitutes condition. Nonetheless, the uniform-price auction is susceptible to demand reduction — meaning that bidders are likely to reduce their demands and thereby end the auction before reaching a Walrasian equilibrium. Indeed, we know from the Fundamental Theorem of Welfare Economics that the Walrasian equilibrium is efficient, so that any conclusion of inefficiency in a uniform-price auction implies that the outcome must be non-Walrasian.

C. Static Pay-as-Bid Combinatorial Auctions

Let us consider an example with two bidders, 1 and 2, and two items, A and B, where the substitutes condition is not satisfied and the existence of Walrasian equilibrium fails. Bidder 1 has a valuation of 3 for the package of A and B, but has a valuation of 0 for each item separately. (Thus, for Bidder 1, the goods are complements — not substitutes.) Bidder 2 has a valuation of 2 for item A, 2 for item B, and only 2 for the package of A and B. The efficient allocation assigns both items to Bidder 1. Consequently, any Walrasian equilibrium (if it exists) must assign both items to Bidder 1. However, to dissuade Bidder 2 from purchasing either item, the prices p_A and p_B of items A and B, respectively must satisfy $p_A > 2$ and $p_B > 2$. Consequently, $p_A + p_B > 4$, exceeding Bidder 1's valuation for the package of two items and yielding a contradiction.

Given the argument of the previous paragraph, we should not expect the simultaneous ascending auction — or any auction format with bids for individual items — to generate the efficient allocation in this example. Bidder 1's dilemma is often referred to as the *exposure problem*: a bidder may refrain from bidding more than his stand-alone valuations for each of the individual items, knowing that if he is outbid on some of the individual items, he will remain "exposed" as the high bidder on the remaining items. This may prevent the available synergies from being realized. Indeed, if Bidder 1 understands this example, he may be unwilling to bid any positive price for either item, since Bidder 2 is sure to win one of the items and therefore Bidder 1 would obtain zero value from the item that he wins.

The exposure problem can be avoided by using a *combinatorial auction*. The rules are modified to permit bidders to place *package bids*, each comprising a *set* of items and a price. For example, the bid ({A, B}, *p*) is interpreted as an all-or-nothing offer in the amount of *p* for the package of A and B — with no requirement that the bidder is willing to accept a part of the package for a part of the price. The allocation is determined by a combination of compatible bids that maximizes the seller's revenues. In this example, Bidder 2 is unwilling to bid any more than 2 for any combination of items, while Bidder 1 is able to exceed 2 for {A, B}. Consequently, the solution has Bidder 1 receiving both items, the efficient allocation.

To the extent that bidders value some of the items in the auction as substitutes, then it may be important for any two bids by the same bidder to be treated as *mutually exclusive*. For example, Bidder 2 in the above example may have been willing to bid 1.5 for item A and 1.5 for item B — but *not* if there was a significant risk that both bids would be accepted. This difficulty is avoided if the auction rules permit at most one of his bids to be accepted. (Such mutually exclusive bids are sometimes referred to as "XOR" bids.) Observe that a rule of mutual exclusivity is *fully expressive* in the sense that it enables the bidder to express any arbitrary preferences. For example, if Bidder 2 in the above example wished to allow both of his bids to be accepted, he could effectively opt out of the mutual exclusivity by submitting a third bid comprising the package {A, B} at a price of 3.

In a static pay-as-bid combinatorial auction, each bidder simultaneously and independently submits a collection of package bids. The auctioneer then solves the *winner determination problem*: find a combination of bids (at most one from each bidder) that maximizes the seller's revenues subject to the constraint that each item can be allocated to at most one bidder. The submitter of each bid selected in the winner determination problem wins the items specified in the bid and pays the amount of the bid.

Rassenti, Smith and Bulfin (1982) are credited with the first experimental study of combinatorial auctions. They studied a static combinatorial auction treating the problem of allocating airport time slots, a natural application given that landing and takeoff slots are strong complements. Bernheim and Whinston (1986) provided an important characterization of equilibria of static pay-as-bid combinatorial auctions under complete information.

D. The Vickrey-Clarke-Groves (VCG) Mechanism

Just as the payment rule of a pay-as-bid auction for a single item or for homogeneous goods can be modified to be "second-price", an analogous modification can be done in the case of a combinatorial auction for heterogeneous goods. This generalization is due to Clarke (1971) and

Groves (1973). Let *N* be an arbitrary finite set of items and let *L* be the set of bidders. In the *Vickrey-Clarke-Groves (VCG) mechanism*, each bidder $\ell \in L$ submits $2^{|N|}$ package bids, for all subsets of set *N*. After the bids are submitted, the auctioneer finds a solution, $(x_\ell)_{\ell \in L}$, to the winner determination problem. While bidder ℓ is allocated the subset $x_\ell \subset N$, he does not pay his bid $b_\ell(x_\ell)$. Rather, his payment $y_\ell \in \mathbb{R}$ is calculated so that $b_\ell(x_\ell) - y_\ell = R^*(L) - R^*(L \setminus \ell)$, where $R^*(L)$ denotes the maximized revenue of the winner determination problem with bidder ℓ present and $R^*(L \setminus \ell)$ denotes the maximized revenue of the winner determination problem with bidder ℓ absent. With sincere bidding, each bid $b_\ell(x_\ell)$ corresponds to the bidder's valuation $v_\ell(x_\ell)$, and $R^*(L)$ corresponds to the (maximized) social surplus. Thus, bidder ℓ is allowed a payoff equaling the *incremental surplus* that he brings to the auction. As in the Vickrey auction for homogeneous goods, a bidder's payment thus equals the opportunity cost of assigning the items to the bidder.

Applied to a setting with a single item, observe that the VCG mechanism reduces to the sealed-bid second-price auction. Applied to a setting of homogeneous goods and nonincreasing marginal valuations, the VCG mechanism reduces to the (multi-unit) Vickrey auction. By the same reasoning as before, the dominance properties of these special cases extend to the setting with heterogeneous items: if bidders have pure private values, sincere bidding is a weakly dominant strategy for every bidder, yielding an efficient allocation.

E. Dynamic Combinatorial Auctions

In auctions for a single item, we have seen that a close relationship exists between a dynamic procedure with a pay-as-bid payment rule (i.e., the English auction) and a static procedure with a second price rule (i.e., the sealed-bid second-price auction). Furthermore, for homogeneous goods with nonincreasing marginal values, an analogous relationship holds between the dynamic Ausubel auction and the static Vickrey auction. An important question for heterogeneous goods is the extent to which outcomes of a dynamic combinatorial auction with a pay-as-bid rule map to the static VCG mechanism.

Banks, Ledyard and Porter (1989) conducted an early and influential study of dynamic combinatorial auctions. They defined several alternative sets of rules for the auction, developing some theoretical results and conducting an experimental study. Other important contributions have included Parkes and Ungar (2000), who independently provided a formulation of the ascending proxy auction described below, and Kwasnica, Ledyard, Porter and DeMartini (2005).

Ausubel and Milgrom (2002) give two formulations of a combinatorial auction and use them to provide a partial answer to the relationship between dynamic combinatorial auctions and the VCG mechanism:

ASCENDING PACKAGE AUCTION. Bidders submit package bids in a sequence of bidding rounds. Each new bid must exceed the bidder's prior bids for the same package by at least a minimum bid increment. After each round, the winner determination problem is solved, on all past and present bids, to determine a provisional allocation and provisional payments. The auction concludes after a round in which no new bids are submitted.

ASCENDING PROXY AUCTION. Each bidder enters his valuations for the various packages into a *proxy bidder*. The proxy bidders then bid on behalf of the bidders in an ascending package auction in which the minimum bid increment is taken arbitrarily close to zero.

The second formulation may be viewed both as a new auction format which greatly speeds the progress of the auction, as well as a modeling device for obtaining results about the first formulation. While the first formulation is an extremely complicated dynamic game, efficiency results and a partial equilibrium characterization are available for the second formulation.

A bidder ℓ in the ascending proxy auction is said to bid *sincerely* if he submits his true valuation, $v_\ell(S)$, for every package $S \subset N$; and he is said to bid *semi-sincerely* if he submits his true valuation less a positive constant, $v_\ell(S) - c$, where the same constant c is used for all packages S with valuations of at least c. The following results refer to the coalitional form game (with transferable utility) corresponding to the package economy: the value of any coalition that includes the seller is the total value associated with an efficient allocation among the buyers in the coalition; and the value of any coalition without the seller equals zero. The *core* is defined as the set of all payoff allocations that are feasible and upon which no coalition of players can improve.

Ausubel and Milgrom (2002) establish that the payoff allocation from the ascending proxy auction, given any reported preferences, is an element of the core (relative to the reported preferences). Furthermore, for any payoff vector π that is a bidder-Pareto-optimal point in the core, there exists a Nash equilibrium of the ascending proxy auction with associated payoff vector π. Conversely, for any Nash equilibrium in semi-sincere strategies at which losing bidders bid sincerely, the associated payoff vector is a bidder-Pareto-optimal point in the core.

Furthermore, the set of all economic environments essentially dichotomizes into two cases. First, if all bidders' preferences satisfy the substitutes condition, then a single point in the core dominates all other points in the core for every bidder, and it equals the payoff vector from the Vickrey-Clarke-Groves mechanism. Thus, in this first case, the outcome of the ascending proxy auction coincides with the outcome of the VCG mechanism. Second, if at least one bidder's preferences violate the substitutes condition, then there exists an additive preference profile for the remaining bidders such that there is more than one bidder-Pareto-optimal point in the core. In this second case, the VCG payoff vector is *not* an element of the core; and the low revenues of the VCG mechanism may become problematic.

7. Conclusion

The proportion of goods and services transacted by auction processes has dramatically increased in recent years and is likely to increase further, making the understanding of auctions and the improvement of their designs increasingly important. At the same time, auctions will remain one of the most useful test-beds for game theory, since the rules of the game are better defined than in most other markets. Consequently, auction theory will almost certainly continue to be a central area of study in economics.

Bibliography

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*, 96(3): 602–629.

Ausubel, Lawrence M. and Peter Cramton (2002), "Demand Reduction and Inefficiency in Multi-Unit Auctions," University of Maryland Working Paper 96–07, revised July 2002.

Ausubel, Lawrence M. and Paul R. Milgrom (2002), "Ascending Auctions with Package Bidding," *Frontiers of Theoretical Economics*, 1: 1–45, http://www.bepress.com/bejte/frontiers/vol1/iss1/art1.

Ausubel, Lawrence M. and Jesse Schwartz (1999), "The Ascending Auction Paradox," mimeo, University of Maryland.

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

Bernheim, B. Douglas and Michael Whinston (1986), "Menu Auctions, Resource Allocation and Economic Influence," *Quarterly Journal of Economics*, 101(1): 1–31.

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

Cramton, Peter, Yoav Shoham and Richard Steinberg (2006), *Combinatorial Auctions*, Cambridge, Massachusetts: MIT Press.

Crémer, Jacques and Richard P. McLean (1985), "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist When Demands Are Interdependent," *Econometrica*, 53(2): 345–61.

Groves, Theodore (1973), "Incentives in Teams," *Econometrica*, 41(4): 617–631.

Gul, Faruk and Ennio Stacchetti (1999), "Walrasian Equilibrium with Gross Substitutes," *Journal of Economic Theory*, 87(1): 95–124.

Harris, Milton and Artur Raviv (1981), "Allocation Mechanisms and the Design of Auctions," *Econometrica*, 49(6): 1477–1499.

Kelso, Alexander S., Jr. and Vincent P. Crawford (1982), "Job Matching, Coalition Formation, and Gross Substitutes," *Econometrica*, 50(6): 1483–1504.

Klemperer, Paul (2000), *The Economic Theory of Auctions*, Cheltenham, UK: Edward Elgar.

Krishna, Vijay (2002), *Auction Theory*, San Diego: Academic Press.

Kwasnica, Anthony M., John O. Ledyard, David Porter and Christine DeMartini (2005), "A New and Improved Design for Multiobject Iterative Auctions," *Management Science*, 51(3): 419–434.

Maskin, Eric and John Riley (1989), "Optimal Multi-unit Auctions," in Frank Hahn (ed.), *The Economics of Missing Markets, Information, and Games*, Oxford: Oxford Univ. Press, 312–335.

Maskin, Eric and John Riley (2003), "Uniqueness of Equilibrium in Sealed High-Bid Auctions," *Games and Economic Behavior*, 45(2): 395–409.

McAfee, R. Preston and John McMillan (1997), "Auctions and Bidding," *Journal of Economic Literature*, 25(2): 699–738.

Milgrom, Paul (2000), "Putting Auction 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 Univ. Press.

Milgrom, Paul R. and Robert J. Weber (1982), "A Theory of Auctions and Competitive Bidding," *Econometrica*, 50(5): 1089–1122.

Myerson, Roger B. (1981), "Optimal Auction Design," *Mathematics of Operations Research*, 6(1): 58–73.

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

Rassenti, S.J., V.L. Smith, and R.L. Bulfin (1982), "A Combinatorial Auction Mechanism for Airport Time Slot Allocation," *Bell Journal of Economics*, 13(2): 402–417.

Riley, John G. and William F. Samuelson (1981), "Optimal Auctions," *American Economic Review*, 71(3): 381–392.

Vickrey, William (1961), "Counterspeculation, Auctions and Competitive Sealed Tenders," *Journal of Finance*, 16(1): 8–37.

Wilson, Robert (1979), "Auctions of Shares," *Quarterly Journal of Economics*, 93(4): 675–689.

Wilson, Robert (1992), "Strategic Analysis of Auctions," in R. Aumann and S. Hart (eds.), *Handbook of Game Theory*, Vol. 1, Ch. 8, Amsterdam: North-Holland.