Simplified Mechanisms with Applications to Sponsored Search and Package Auctions

Paul Milgrom[1]
First draft: August 16, 2007
This draft: October 13, 2007

A simplified mechanism is a direct mechanism modified by restricting the set of reports or bids. An example is the auction used to place ads on Internet search pages, in which each advertiser bids a single price to determine the allocation of eight or more ad positions on a page. If a simplified mechanism satisfies the "best-reply-closure" property, then all Nash equilibria of the simplified mechanism are also equilibria of the original direct mechanism. For search advertising auctions, suitable simplifications eliminate inefficient, low-revenue equilibria that are favored in the original direct mechanism when bidding costs are positive.

Keywords: simplified mechanisms, sponsored search, generalized second price auctions, package auctions, combinatorial auctions, core-selecting auctions, mechanism design.

JEL Categories*:* D44, C78

I. Introduction

Real-world auctions are often much simpler than the direct mechanisms which are the focus of economic theory. The latter are impractical for large auctions with many items for sale, for the sheer number of combinations of items makes determining and

[1] Support for this research was provided by National Science Foundation grant SES-0648293 and by Yahoo! Thanks to Richard Steinberg for comments. Any opinions expressed here are those of the author alone.

expressing all the relevant values too costly. Simplified mechanisms that reduce the number of required bids can help, but raises the question of how the mechanism's performance is affected. This paper initiates the study of that question with the finding that well-chosen simplifications can sometimes improve mechanism performance by eliminating undesirable equilibria.

One way to reduce the number of bids used by a direct mechanism is to conflate distinct allocations or events, so that a bid that applies to one is required to apply also to others. A very common example arises in auctions where each bidder is allowed only to specify values or bids that depend only its own goods assignment, rather than on the entire allocation. This example and certain others seem easily explained as respecting an approximate pattern of preferences, so it is more enlightening and convincing to look at an example in which conflation is used in a different way.

An interesting simplified auction mechanism is the world's most frequently used auction, which is initiated whenever a user types text into in a search box like those provided by Yahoo! or Google. For each search, an automated auction runs to determine the placement of advertisements into multiple positions—currently eight at Google and twelve at Yahoo!—on the search results page. In the earliest search auctions, bidders in an automated auction could offer a separate price for each position and a sequence of first-price auctions determined the winners. Each bidder in the series of auctions was permitted to win a single ad position.[2] In this early incarnation, the auction entailed no simplification.

[2] Search auctions were pioneered by Overture in 1998. The company was later purchased by Yahoo! Google introduced its search advertising auction in 2000.

As currently administered, the would-be participants in a search auction identify search terms that will trigger their bids and specify a price per click for each term; they also specify the ad to run if the bid is winning. The auctioneer converts the per-click bid of each bidder into a per-impression bid by multiplying it by an estimate of the number of clicks the ad would experience if it were shown in the first position on the page. The highest bidder wins the top position; the second highest wins the next position; and so on. An advertiser pays only when its ad is actually clicked and then pays only the smallest bid per click that would have won the same ad position.

This mechanism, dubbed the "generalized second-price auction," is equivalent to a series of second price auctions with separate per-impression bids for each position, but with two restrictions. The first restriction is similar to that of earlier search advertising auctions: an advertiser who wins one position on a page is excluded from bidding for the lower positions. The second restriction is a simplification: an advertiser's per-impression bid for the n^{th} position in the sequence is determined by its bid for first position, but scaled down in proportion to the lower number of expected clicks for the n^{th} position.

In a pair of recent papers, Edelman, Ostrovsky, and Schwartz (2007) and Varian (2006) have studied the generalized second-price auctions using the assumptions that bidders value all clicks on ads equally (regardless of the position of the ad) and that bidder payoffs are equal to the value of their clicks minus the total prices they pay. A central finding of both papers is that the prices and assignments of positions resulting from a selected full-information Nash equilibrium of the generalized second price auction is the same as for the dominant strategy equilibrium of a multi-item Vickrey auction.

This theory leaves several unanswered questions about sponsored search auctions. First, why do advertisers pay on a per-click basis, rather than on the per-impression basis that is most commonly used for print ads and for radio and television advertising? In a static full-information environment, there would be little to distinguish between these different approaches to pricing, although per-click charges are easier for an Internet advertiser to audit because it can meter visits to its own website.

There is a second important advantage to per-click pricing. Search companies have continually expanded their scope in various ways, showing ads on a wider variety of sites and encouraging advertisers to use "extended match" technologies to place ads not only on pages that match the bidder's search term exactly but also on pages that match approximately. As an illustration, the extended search technology might deem the term "ink cartridge" to be sufficiently related to the term "printer cartridge" and might show an ad for the latter when the search is made for the former. The relation among these search terms is imperfect, for example because "ink cartridge" might be entered by a user searching for a pen ink refill, so the proportion of searchers who are potential customers for a printer ink company may be lower for the related term, which makes each impression less valuable. Even click values may be different, because clicks from pen ink searchers would less frequently result in actual sales. Still, pricing ads on a per click basis reduces the advertiser's cost per impression for ads on less closely related search results pages, which makes more advertisers willing to agree to use the extended search technology. This explanation is part of a recurring theme of our analysis: per-click bidding is a simplification that reduces the number of bids required and increases the scope of each, raising reported demand and increasing the seller's revenues and profits.

Another question concerns not the distinction between price-per-click and price-per-impression bids, but the choice of auction rules. If, as the prior literature asserts, the Vickrey outcome is a desirable one, then why not just use the Vickrey mechanism instead of the generalized second-price auction?[3] Not only does that mechanism implement the desired outcome using dominant strategies rather than merely full-information Nash equilibrium, but it does so for a realistically wider class of environments in which the value of a click may depend on the position of the ad in addition to the search term.

We divide this relatively broad question into two narrower ones by treating the generalized second-price auction as differing from the Vickrey mechanism in two ways. First, its bids are simplifications: there are one-dimensional while the value reports required by the Vickrey mechanism (and those required by earlier search auctions) are multidimensional. Second, given the vector of values that might be imputed from the one-dimensional bids, the pricing of the ad positions is determined not by the Vickrey formula but by a sequence of second price auctions.[4] These two differences suggest two corresponding questions. First, if pricing is to be set by a sequence of second price auctions, why does the auctioneer then accept only a single price per click and impute values to all positions instead of allowing multidimensional bids that state directly all the relevant values? Would the same explanation apply if the Vickrey pricing rule had been used? Second, if the auctioneer must use a single, one-dimensional bid and impute values

[3] In early postings describing the auction, Google claimed that this generalized second price auction was the actual Vickrey auction, but that is a mistake. In particular, no bidder has dominant strategy in the generalized second price auction.

[4] Although this represents just one particular way to represent the generalized second price auction as a simplified direct mechanism, this representation is a useful benchmark because a series of single item auctions is the most common way to sell similar heterogeneous items in many auction settings.

for the different positions, what advantage might it enjoy by using a sequence of second-price auctions rather than the Vickrey auction?

To answer the first question, we observe that in any series of second-price auctions, it is the losing bids for the various positions that determine the prices. If individual bids for each position were permitted but not required and if there were any arbitrarily small positive cost incurred by a bidder in submitting individual bids, then there would be no pure, full-information equilibrium at which the seller earns positive revenue, because losing bidders for a position would never make positive bids.[5] Even when the cost of submitting bids is zero, the series of second price auctions with individual bids still admits these zero-revenue strategy profiles as Nash equilibria. Similar arguments imply that the Vickrey pricing rule never yields positive revenue in a pure, full-information equilibrium when there are positive bidding costs and that these zero-revenue equilibria persist even when bidding costs are zero. In contrast, every equilibrium of the generalized second price auction for two or more items awards positive revenues to the seller, because a bidder whose positive bid is winning for position n also enforces a positive price for position n–1. We will argue below that this analysis, which here seems tailored to exploit the particular structure of the generalized second price and Vickrey auctions, nevertheless applies more broadly and illustrates a general principle of mechanism design: certain kinds of simplifications *reduce* the set of pure Nash equilibria—often by eliminating inefficient or low-revenue equilibria—*without* introducing additional pure equilibria.

[5] This paper uses full-information Nash equilibrium to analyze various mechanisms. Based on earlier empirical successes and failures of game-theoretic auction models, what we believe should be taken most seriously from this analysis is the *comparative* predictions about the revenue performance of alternative auctions mechanisms, rather than the point predictions about the performance of any single mechanism.

For the second question, although the received literature already includes analyses highlighting important disadvantages of the Vickrey pricing formula in multi-item auctions (Ausubel and Milgrom (2005), Rothkopf (2007)), the most devastating objections apply only to auctions in which bidders can buy multiple items. Those objections have no force for sponsored search auctions, because each bidder in such an auction is restricted to buy at most a single position.

Our answer to the second question focuses on the special environment of sponsored search, for which a distinct analysis is needed. We extend the models used in earlier studies to allow heterogeneity among searchers. We assume that there are two kinds of searchers—some are potential customers who are actually looking for a product to buy and others are merely curious about the products being advertised—with each group having its own click rates for ads occupying different positions on the page. For example, it may be that clicks on ads near the bottom of a search page come more frequently from potential customers because these searchers more often attend to the full list of the ads. In that case, if clicks from potential customers are more valuable than clicks from other searchers, then clicks on ads near the bottom of a page will be more valuable than clicks on ads near the top, because a higher proportion of these clicks will come from potential buyers. In general, we need only assume that the click rates for the groups are different to conclude that clicks from different positions have different values.

Our formal model incorporates searcher heterogeneity in a simple way by assuming that there are just two groups of searchers: potential customers and others. Each advertiser has some positive value per click from potential customers and a zero value per click from other searchers and the frequency of clicks from each group falls as one moves down the search page. With these assumptions, the bidders' types are one-dimensional and the value per impression falls as one moves down the page, just as in the prior literature. Based on the data at its own site, the auctioneer can observe the empirical click rate for each position but not the purchase behavior of clickers once they leave the search page. The auctioneer cannot determine from its own observations and a bidder's reported value for an ad in one position what the bidders' values are for ads in the other positions. Therefore, with one-dimensional bids, it has too little information to conduct a proper Vickrey auction despite the one-dimensional type spaces. In contrast, the analyses of the previously cited papers can be generalized to establish that, even with searcher heterogeneity, there may still exist a full-information equilibrium of the generalized second price auction in which the realized prices are Vickrey prices. This is possible because each bidder can observe how its own clicks from various ad positions convert into sales and profits.

The lessons illustrated about the advantages of limited bidding in sponsored search auctions suggest a more expansive theory of simplified mechanisms, which are derived from direct mechanisms by restricting the set of allowable reports or bids. A key characteristic of successful simplifications is the best-reply-closure property, defined as follows: for any participant *j*, if the other participants play only their own simplified strategies, then participant *j*'s set of simplified strategies includes a best reply to the profile of others' strategies. The simplification used in sponsored search auctions, in which each bidder names a single price rather than a vector of prices, satisfies the best-reply closure property. Our main general theorem asserts if a simplified mechanism has the best-reply closure property, then a profile of pure, simplified strategies is a Nash equilibrium of a simplified mechanism if and only if it is a Nash equilibrium of the original mechanism. Such a simplification can eliminate pure equilibria (by eliminating one or more of the strategies it uses) but otherwise leaves the set of equilibria unchanged.

Besides Internet search advertising, a second significant application of simplified mechanisms is to the problem of *package auctions* (also known as *combinatorial auctions*). These are mechanisms in which there are multiple (often heterogeneous) items for sale and bidders are potentially interested in buying any packages, that is, subsets of the full set of items. With M items for sale and quasi-linear preferences, a full description of a bidder's preferences specifies values for all $2^M - 1$ non-empty packages. If a direct package auction mechanism were attempted for a sale like FCC spectrum auction #66 in which 1122 licenses were offered for sale, a bidder could feasibly compute and report values for only an extremely minute fraction of the roughly 10^{338} available packages. If we model this fact by assuming that bidders can submit a modest number of packages bid at no cost but eventually incur a small cost for each additional package bid, then there can be a huge number of inefficient and low-revenue equilibria of the full game. We examine how a simplified package auction satisfying the best-reply closure property can eliminate certain "undesirable" equilibria without introducing new Nash equilibria.

Our analysis of simplified package bidding treats the class of core-selecting package auction mechanisms of Day and Milgrom (2007)—a class of direct mechanisms that includes the important menu auction of Bernheim and Whinston (1986), the ascending proxy auction of Ausubel and Milgrom (2002), and many others. For these mechanisms, the full-information equilibrium outcomes include all the bidder-optimal core allocations.

One creates a simplified mechanism from a direct mechanism by restricting bidders to report values that are elements of a set *V*. With a set of items *N* for sale, a typical element $v \in V$ is a function $v: 2^N \to \mathbb{R}_+$ with the property that $v(\emptyset) = 0$. **For $k > 0$, let $v - k$ denote the value function which assigns to any non-empty package *S* the value $v(S) - k$. We show that if the actual values lie in the set *V* and if $v \in V \Rightarrow v - k \in V$, then the best-reply-closure property is satisfied. Consequently, the Nash equilibria of the *V*-simplified mechanism are Nash equilibria of the original mechanism, and these include the identified equilibria for which the outcomes are bidder-optimal core allocations.

Based on the preceding analysis, we suggest some sets *V* that may be useful for applications in which potential value complementarities arise only from shared fixed costs. One useful property of our sets *V* is that they grow only linearly in the number of items *N*, while the full set of package bids grows exponentially in *N*. We evaluate the performance of these simplified mechanisms in particular environments, including ones in which the actual values lie outside of *V*. This analysis allows us to revisit the difficult question of whether, when and how prices might be useful in package auction design.

The rest of this paper is organized as follows. Section II states and proves the *simplification theorem*, which shows that for general games, simplifications that restrict the strategy set to one satisfying the best-reply-closure property shrinks the set of pure Nash equilibrium profiles. We also identify a common property of standard auctions—*completeness* of the set of bids—that is sufficient to imply the best-reply closure property. Section III treats the generalized second price auction of sponsored search.[6] Its

[6] Throughout our analysis of auctions, we set aside the possibility of ties. These can be treated by an extension of the equilibrium concept, as suggested by Simon and Zame (1990), or by other devices, but these details do not affect any substantive conclusions.

first subsections shows that, compared to a series of second price auctions with general value reports, the generalized second price auction is a *complete* simplification (so it satisfies the best-reply-closure property) and eliminates certain zero revenue Nash equilibria. Its second subsection introduces the model described above with two types of searchers and demonstrates that the selected equilibrium of the generalized second price auction still establishes Vickrey prices, thus extending the results of prior research. Section IV treats package bidding, proving the theorem stated above which identifies a class of simplifications that satisfies the best-reply-closure property. Section V concludes.

II. Simplification and Completeness Theorems

Let (N, X, π) be a normal form game, where $X = (X_1, ..., X_N)$.

Definition. A product set of strategy profiles $\hat{X} = \hat{X}_1 \times ... \times \hat{X}_N$ has the *best-reply closure property* in (N, X, π) if for every player *n* and every profile $x_{-n} \in \hat{X}_{-n}$ there exists $x_n \in \hat{X}_n$ such that for all $x_n' \in X_n$. $\pi_n(x_n, \hat{x}_{-n}) \geq \pi_n(x_n', \hat{x}_{-n})$.

When the best-reply closure property holds, a player *n* looking for a response to any opposing pure profile $x_{-n} \in \hat{X}_{-n}$ loses nothing by restricting attention to strategies in $\hat{X}_n$.

Theorem 1 (Simplification Theorem). Suppose $\hat{X}$ has the best-reply closure property in (N, X, π). Then, a pure strategy profile $\hat{x} \in \hat{X}$ is a Nash equilibrium of $(N, \hat{X}, \pi)$ *if and only if* it is also a Nash equilibrium of (N, X, π).

Proof. The *if* direction is obvious. For the only if direction, suppose that $\hat{x}$ is not a Nash equilibrium of (N, X, π). Then there is some player *n* that has a profitable deviation from $\hat{x}$, that is, for some $x'_n \in X_n$, $\pi_n(x'_n, \hat{x}_{-n}) > \pi_n(\hat{x}_n, \hat{x}_{-n})$. According to the best-reply closure property, there is some $x_n \in \hat{X}_n$ such that $\pi_n(x_n, \hat{x}_{-n}) \geq \pi_n(x'_n, \hat{x}_{-n})$. Hence, $\pi_n(x_n, \hat{x}_{-n}) > \pi_n(\hat{x}_n, \hat{x}_{-n})$: $\hat{x}$ is not a Nash equilibrium of $(N, \hat{X}, \pi)$. ♦

The interesting part of the simplification theorem is the *only if* assertion. It says that eliminating strategies while preserving the best-reply closure property does not add new equilibrium strategy profiles and hence does not extend the set of equilibrium outcomes. For applications, the trick is to specify $\hat{X}$ to eliminate the "bad" equilibria while preserving the "good" equilibria and to verify the property, so that no new bad equilibria are introduced.

The simplification theorem has been stated above for equilibria in pure strategies and we will apply it in that form. Since mixed strategy equilibria are pure equilibria of a game with an enlarged strategy space, there is a corollary for the mixed equilibrium case, but it uses the stronger *mixed-best-reply closure condition*. We state that condition as follows: for every mixed strategy profile $\delta_{-n} \in \times_{j \neq n} \Delta(\hat{X}_j)$, there exists $x_n \in \hat{X}_n$ such that for all $x'_n \in X_n$. $\pi_n(x_n, \hat{x}_{-n}) \geq \pi_n(x'_n, \hat{x}_{-n})$.

Theorem 2. Suppose $\hat{X}$ has the mixed-best-reply closure property in (N, X, π). Then, a profile $\delta \in \times_j \Delta(\hat{X}_j)$ is a mixed Nash equilibrium of $(N, \hat{X}, \pi)$ *if and only if* it is also a mixed Nash equilibrium of (N, X, π).

The best-reply closure property is useful because it is satisfied by many typical simplifications of direct multi-item auction mechanisms. One condition that implies it for auctions is based on the how the set of reports is restricted. For this development, we distinguish between reports, which specify a whole set of values, and "bids," which specify the value of a particular set. A "winning bid" is then be the value a bidder reports for the particular item or package of items assigned to it; all other bids are "losing bids."

Our main result depends on two definitions.

Definition. An auction mechanism is *standard* if (1) it is a (possibly simplified) direct mechanism and (2) the payment required from any bidder is a function of the reports of the other participants and of its own winning bid (and does not depend on the bidder's own losing bids).

Definition. A simplification of an auction mechanism is *complete* when for each bidder j, each report by the other bidders v_{-j}, each package x_j, and each price p_j, if there is some report v_j in the underlying mechanism that wins package x_j with winning bid p_j, then there is a report in the restricted set that wins package x_j with winning bid p_j.

From these two definitions, the following is immediate.

Theorem 3 (Completeness). A complete, standard auction mechanism satisfies the best-reply closure property.

Completeness is a restriction on the way the set of bids is reduced when an auction is simplified. Standard auctions are derived from direct mechanisms and encompass all the direct mechanisms commonly discussed in the economics literature.

III. Application to Search Auctions

For this section, we follow the earlier literature by treating bids as prices per impression rather than prices per click. As we have already described, this conversion is straightforward when search terms are interpreted narrowly; it does not affect the strategic analysis in that case.

Simplified Search Auctions Are Desirable

Suppose that bidder *i*'s value of an ad in position *n* is denoted v_{in}. Each advertiser is permitted to acquire only one ad position, so the vector v_i completely describes the bidder's values for the possible positions it might acquire. We make the standard normalization that a bidder who gets no ad has a zero payoff. Let us initially suppose that there is a small cost ε of submitting a positive bid for each position. In this model, there is no best reply to any pure strategy profile that entails a positive losing bid, so in particular the usual dominant strategy analysis for the Vickrey auction fails. That analysis does, however, have a useful counterpart in the model with costly bidding: if bidder *i* submits a positive bid $v_{in} \neq b_{in} > 0$ for just one position, then that bid is weakly dominated by $b_{in} = v_{in}$. By inspection, if bidders bid only for the items that would be assigned to them in an efficient allocation, then the corresponding singleton bids $b_{in} = v_{in}$ describe a Nash equilibrium. Summarizing:

Theorem 4. In any pure strategy equilibrium of the Vickrey auction game with costly bidding, the seller's revenue is zero. If the equilibrium bids are undominated, then the winner *i* of position *n* bids $b_{in} = v_{in}$ for that position. There is a zero-revenue undominated equilibrium in which the items are assigned efficiently. This efficient zero-revenue equilibrium bid profile is also a (dominated) pure Nash equilibrium when the bid cost is zero.

The Vickrey auction thus has undesirable Nash equilibrium properties when there is even an arbitrarily small cost of reporting bids. To make an analogous statement for a series of second price auction, we let the vector $b_i = (b_{i1},...,b_{iN})$ denote the bids that advertiser *i* is prepared to make for each of the *N* positions. To keep notation simple, let us permute the bidder indexes so that bidder 1 is the bidder who wins the first position, bidder 2 the second, and so on. Let $L_n = \max_{j>n} b_{jn}$ denote the second highest ("losing") bid for position *n*. In the sequence of second-price auctions, this is the price paid by bidder *n* to acquire ad position *n*. If bidder *n* makes K_n positive bids, then its payoff is $v_{nn} - L_n - \varepsilon K_n$.

<u>Theorem 5</u>. In any pure strategy equilibrium of the sequence of second price auctions with costly bidding, the seller's revenue is zero. If the equilibrium bids are undominated, then the winner *i* of position *n* bids $b_{in} = v_{in}$ for that position. There is a zero-revenue undominated equilibrium in which the items are assigned efficiently. This same bid profile is also a pure Nash equilibrium when the bid cost is zero.

In both the Vickrey auction and the sequence of second price auctions, the revenue result reverses when the strategy sets are simplified.

For the Vickrey auction, suppose we follow the earlier papers in assuming that bidder values per click do not depend on the ad position and that the click rate on an ad in position *n* is some fixed from α_n of the rate in position 1, where $1 = \alpha_1 > ... > \alpha_N > 0$. Then, $v_i = v_{i1}(1, \alpha_2, ..., \alpha_N)$; the bidder's value space is one-dimensional. The auctioneer needs only to ask each bidder for a bid b_{i1} for the first position. Since the auctioneer can observe α, it can compute the Vickrey prices for each bidder and position. In the resulting game, if there are positive bidding costs, any bid $v_{i1} \neq b_{i1} > 0$ is weakly dominated by the bid $b_{i1} = v_{i1}$ .In an undominated pure equilibrium, each of the bidders with the *N* highest values will prefer to make positive bids and the other bidders will prefer to bid zero. Position *N* will have a price of zero, but the price of any position $n < N$ is at least $(\alpha_n - \alpha_N) b_{N1} > 0$, since the opportunity cost of position *n* is not less than the gain from reassigning bidder *N* to that more valuable position.

Theorem 6. With $N > 1$ positions for sale, at least *N* bidders, and zero or small positive bidding costs, there is no zero-revenue equilibrium of the simplified Vickrey auction. At any pure equilibrium, the price paid for position *N* will be zero, but all other prices will be strictly positive.

A similar analysis applies to using single bids for a sequence of second-price auction. This is precisely the generalized second-price auction.

Theorem 7. With *N*>1 positions for sale and zero or small positive bidding costs, there is no zero-revenue equilibrium of the generalized second price auction. The price paid for position *N* will be zero, but all other prices will be strictly positive.

Only the cases with zero bidding costs are formally applications of the Simplification Theorem. For those cases, the zero-revenue Nash equilibria are eliminated by simplifying the strategy set for the Vickrey auction or the series of second-price auctions, but certain positive revenue equilibria remain. We have included positive bidding costs in this analysis because they select certain interesting equilibria and because they are an integral part of the reason for making simplifications, providing a bridge connecting the theories of sponsored search and package bidding.

The One-Dimensional Vickrey Pricing Rule is Undesirable

We have just seen that, in a particular model, a simplification that enables the auctioneer to implement Vickrey pricing from one-dimensional bids. If Vickrey pricing is both implementable and desirable, why does the search auctioneer not do that? Does to the generalized second-price auction have a heretofore unrecognized advantage?

The answer offered here uses the fact that the preceding analysis incorporates an unjustified assumption, namely, that the value of clicks is independent of the position of the ad. To explore an alternative, we introduce heterogeneity among searchers, supposing that there are two types. Searchers of one type ("potential buyers) are looking for a product to buy while those of the other ("curious searchers") are merely looking for information. The ratio of curious searchers to potential buyers is denoted by λ.

In the prior literature, it is supposed that a searcher's click rate on an ad is determined by multiplying the ad's "clickability" times the click rate for the position. Here, we assume the same. For potential buyers, the relative click rate on an ad in position *n* is α_n; for curious others, it is β_n. We assume that $\alpha_1 > ... > \alpha_N > 0$ and $\beta_1 > ... > \beta_N > 0$, but we do not assume that the two series are proportional. For example, if the attention of curious searchers flags more quickly than that of potential buyers, then the sequence β_n / α_n would be decreasing.

We assume that only clicks by potential buyers are valuable to advertisers, so the value of an ad in position *n* is $v_i \alpha_n$. A bidder can learn this positional value over time by observing its sales from ads in position *n*. The formulation $v_i\alpha_n$ for the matching value implies that assortative matching is efficient, that is, the advertiser with the highest value v_i should be shown in first position, and so on for the other positions. It simplifies the exposition to label the bidders so that $v_1 > ... > v_M$ and to assume that there are weakly more positions than bidders $M \geq N$. Then, at the efficient allocation, position *n* is assigned to bidder *n*.

It has long been known that market clearing prices exist for a class of matching problems including the one described and further that there is a unique minimal market clearing price vector p which can be computed using linear programming (Koopmans and Beckmann (1957)). The minimum equilibrium price p_n is the shadow price of an additional impression in position *n*. It follows that p_n is the opportunity cost of the ad placed in position *n* by bidder *n*, so it is also the Vickrey price paid by bidder *n* to acquire that position.

Competitive equilibrium prices satisfy constraints that bidder *n* prefers position *n* to position *n*–1, that is, $v_n\alpha_n - p_n \geq v_n\alpha_{n-1} - p_{n-1}$ and, as is familiar from mechanism design analyses, the single crossing structure of preferences assumed here ensures that these hold as equalities at the minimum competitive equilibrium. Treating $\alpha_{N+1} = 0 = p_{N+1}$, it follows that the Vickrey prices are $p_n = \sum_{k=n}^{N}(p_k - p_{k+1}) = \sum_{k=n}^{N}\left(v_{k+1}(\alpha_k - \alpha_{k+1})\right)$, which is the formula for such prices reported by Edelman, Ostrovsky, and Schwartz (2007).

The click rate for position *n* is $\alpha_n + \lambda\beta_n$. Although this rate decreases with *n*, it would be a rare coincidence for it to decrease in direct proportion to the value of an ad. Since the search company observes clicks but not sales, it varies bids in proportion to clicks but not in proportion to value. If bidder *i* names a price of b_{i1} for position 1 in a simplified auction, then the auctioneer can impute a bid for position *n* as $b_{i1}\gamma_n$, where $\gamma_n = (\alpha_n + \lambda\beta_n)/(\alpha_1 + \lambda\beta_1)$ is the relative click rate for position *n*, but the auctioneer *cannot* generally infer Vickrey prices from these bids and its other information.

Is the efficient assignment with the Vickrey price vector *p* is the outcome of Nash equilibrium in the generalized second-price auction? If it is, then it must be that the highest bid is made by bidder 1, the second highest by bidder 2, and so on, and that the highest losing bidder for each position bids the Vickrey price for that position. Thus, for each bidder *n* for $n = 2,\ldots, N+1$, it is necessary that the equilibrium bids are $b_{n1} = p_{n-1}/\gamma_{n-1}$. The other bids are not uniquely determined, but we may specify that bidder 1 bids $b_{11} = \alpha_1 v_1$ and that bidders with indexes *N*+1 and larger bid $b_{n1} = \alpha_N v_n / \gamma_N$.

<u>Theorem 8</u>. For the two searcher-type model of this section, there is a pure Nash equilibrium of the generalized second-price auction in which the assignment is efficient and prices paid by the winning bidders are the Vickrey prices *p if and only if* the corresponding price-per-click sequence $\{p_n/\gamma_n\}_{n=1}^{N}$ is decreasing.

<u>Proof</u>. If the Vickrey-price-per-click sequence p_n/γ_n is not decreasing, then the bidders are not ranked in the correct order for an efficient assignment. (For example, if

$p_3 / \gamma_3 < p_4 / \gamma_4$, then bidder 4 bids less than bidder 5 and the resulting assignment is inefficient.)

Suppose that $\{p_n / \gamma_n\}_{n=1}^N$ is decreasing and fix any bidder *n*. Recall that the Vickrey prices are competitive equilibrium prices so no bidder wishes to deviate to purchase a different position at prices *p*. If bidder *n* raises its bid to win a higher position, say position $k < n$, then the price it must pay is determined by the k^{th} highest bidder, so it is $\gamma_k (p_{k-1} / \gamma_{k-1}) > \gamma_k (p_k / \gamma_k) = p_k$, so that deviation is unprofitable. If bidder *n* reduces its bid to win a lower position $k > n$, then the price it must pay is precisely p_k and the deviation is again unprofitable. ♦

Previous literature establishes that the desired equilibrium exists when $\lambda = 0$ or more generally when the vector γ is proportional to the vector α, that is, when the seller's estimate of relative values is not too far off. When the values v_i of the various bidders are very close, then this condition is almost necessary, so the generalized second-price auction does not work well. When the values variation is larger, this constraint is more relaxed.

In any series of second-price auctions in which advertisers other than *j* were obliged to use one-dimensional strategies, suppose that a best reply by *j* wins some position *n*. The price *j* pays in that case is determined by the n^{th} highest opposing bid. It can obtain the same position at the same price with a one-dimensional bid that is the n^{th} highest such bid. Therefore, we have proved the following.

Theorem 9. The generalized second-price auction is a complete, standard auction mechanism.

It follows from theorems 3 and 9 that the generalized second-price auction satisfies the best-reply closure property, so the Simplification Theorem applies: The pure Nash equilibria of the generalized second-price auction are also equilibria of any sequence of second-price auctions with richer strategy sets. As we have seen, there are equilibria of the full direct mechanism that entail zero revenues. The generalized second-price auction as it is actually conducted for sponsored search applications has no such equilibria.

The analysis reported in this section was formulated for application to online search, but similar analyses in which bidders are forced to make the same bids for different items apply to other Internet advertising auctions. What makes this sort of enforced conflation is valuable is that advertising targets can be so highly differentiated. For example, a Palo Alto mortgage lender might be prepared to bid high to target a refinancing online advertisement to "males aged 35-54, homeowners in Palo Alto, CA, with good credit scores whose navigation behavior displays interest in home improvement or mortgage refinance and who are not currently visiting a sex or gambling site." Detailed targeting can be valuable because it improves the matching of ads to users, but narrow targeting can also reduce competition and result in low revenues for online publishers. Sponsored search is just one example in which a simplified auction that conflates distinguishable ad opportunities can increase equilibrium revenues.

IV. Application to Package Auctions

In contrast to the assumption made in much economic theorizing that auctions are conducted for a single item, many auctions take place in settings where multiple items are being sold and the sales interact. This relationship can emerge from budget constraints that prevent independent bidding on separate items. It can also emerge when the goods enter the buyer's production or utility function as substitutes or complements. Although such interactions are very common, *package auctions*, in which bidders can name prices for the packages of lots or items they wish to buy, are only infrequently used.[7] More often, items/lots/tranches are sold sequentially or in simultaneous sealed bids. The use of these alternative arrangements calls for explanation.

It seems intuitively clear that these one-item-at-a-time auctions are simpler than package auctions, although the rubric "simple" is an ambiguous one. One important meaning that has received some attention is that computation is much easier for single item auctions than for package/combinatorial auctions. A second simplicity notion, which we have emphasized in this paper, is that bids are restricted so that bidders are called upon to make fewer bids.[8]

Many common single item auctions are simplified package auctions according to our definition. For example, a simultaneous second-price auction for N items is a simplification of a standard Vickrey package auction for N items in which bidders are allowed to make only bids that express values of packages as the sum of the values of their constituent items. Also, a simultaneous first-price auction is the simplification of a Bernheim-Whinston menu auction with the same bid restriction.

Many more complex package auctions impose restrictions on bids that qualify them as simplified package mechanisms in the sense introduced here. For example, the

---

[7] A recent book by Cramton, Shoham, and Steinberg (2005) reports a snapshot of the growing literature on package auctions, including reports of applications. Milgrom (2004) describes additional applications.

[8] This type of simplicity is relevant for reporting and computation, too, since the amounts of reporting and computing time are functions of the amount of data.

City of London procures bus services using a package auction which requires bidders to submit a price meeting the reserve for each named route while allowing discounts to be offered for combinations of routes (Cantillon and Pesendorfer (2005)).

Below, we limit attention to simplifications of core-selecting package auctions. The underlying direct mechanisms are ones that always select an allocation in the core determined by reported values. Among these mechanisms are the menu auctions studied by Bernheim and Whinston (1986). Those authors showed that for every *bidder-optimal* allocation (meaning a core allocation that is not Pareto dominated for the bidders by any other core allocation), there is a coalition-proof equilibrium of the menu auction which selects that allocation. If π is the corresponding bidder-optimal core imputation, then the equilibrium strategy profile has each bidder *j* report that each non-empty package *S* has value $\max(v_j(S)-\pi_j,0)$, where v_j is the bidder's actual value function for packages. We denote this report by $v_j-\pi_j$.

Day and Milgrom (2007) show that precisely these same profiles of *profit-target strategies* $v_j-\pi_j$ are Nash equilibria of *every* core selecting auction mechanism. They also show that for every core-selecting auction and every strategy profile of the other bidders, bidder *j* has a best reply of the form v_j-k for some $k\geq 0$. The theory we develop below applies to this whole set of auction mechanisms.

Consider a simplified core-selecting auction in which bidders are restricted to report values in a set *V*. With a set of items *N* for sale, a typical element $v\in V$ is a function $v:2^N\to\mathbb{R}_+$ with the property that $v(\varnothing)=0$. For $k>0$, let $v-k$ denote the value function which assigns to any non-empty package *S* the value $v(S)-k$.

Definition. The set of values *V* is *closed under fixed costs* if for all $k > 0$, $v \in V \Rightarrow v - k \in V$.

A direct application of Theorem 2 of Day and Milgrom (2007) yields the following result.

Theorem 10. Let Γ_V be a simplified core-selecting auction with reports restricted to lie in the set *V*. Suppose that *V* is closed under fixed costs and that actual bidder values lie in the set *V*. Then, Γ_V has the best-reply closure property and the profit-target equilibrium strategy profiles identified above for the full mechanism are also equilibrium of the simplified mechanism.

Theorem 10 identifies a class of simplified mechanisms for package bidding. For example, *V* might be the set of values expressed as the sum of item values, minus a constant: $v \in V \Leftrightarrow \left(\exists \alpha \in \mathbb{R}_+^N, k \in \mathbb{R}_+\right)(\forall S \neq \emptyset)\, v(S) = \sum_{n \in S} \alpha_n - k$. Elements of *V* could express values of collections of items when there is a fixed cost of shipping or a shared facility that must be built to use the items. Simplified core-selecting mechanisms using this *V* can be dubbed *fixed cost package auctions*.

Among the important features of the fixed cost package auctions is that they eliminate many (but not all) coordination failure equilibria. For example, suppose that $N = \{1,2,3\}$ and that there are three bidders. Suppose that bidder 1 values only item 1 and has a value of 10; bidder 2 values only items 2 and 3 with values of 10 each and fixed costs of 10, and that bidder 3 values the items at 5 each, with no fixed cost. Among the Nash equilibria of the full menu auction is one at which bidder 3 wins all the items, bidding 15 for the whole set and making no other bids, while bidders 1 and 2 each bid 10 for the whole set, making no other bids. There is no corresponding equilibrium of the simplified game. If bidders 1 and 2 play only undominated strategies and bid their full values for the package of the whole, then the only corresponding equilibrium outcome entails an efficient allocation. This illustrates the Simplification Theorem, according to which the narrower strategy set can eliminate equilibria but cannot introduce additional equilibria.

Two other important advantages of the fixed cost package auction design are the low dimensionality of the reports required from bidders and the fact that for any fixed number of bidders, computation time rises only linearly in the number of items for sale.

Affine Approximation Mechanisms

Here we propose a simplified mechanism that incorporates the fixed cost package auctions while preserving all of its advantages and also extends a design created by the author to sell the generating assets of an electric utility company. In the asset sale application, two kinds of bidders were expected to participate in the auction—ones that wanted to buy all or nearly all of the generating portfolio and others that wanted to buy only specific very small parts of the portfolio. For example, the company's partners in ownership of some electric generating facilities might want to buy the selling company's share in order to avoid being saddled with unfamiliar new partners and counterparties to certain contracts might want to buy back their commitments. The suggested design involved two stages[9] of which the second involved a package auction in which bidders for the whole portfolio of assets would be required to specify decrements to be applied to their bid for the whole portfolio if some of the individual pieces were sold to others.

[9] The first stage involved indicative bids to identify qualified bidders and to determine which assets would be open for individual bidding.

Partners and counterparties bidders could bid for the individual pieces for which they were qualified.

Generally, we define the *affine approximation mechanisms* to be simplified core-selecting auctions in which a bid (T,β,α,r) comprises a package *T*, an offer β for that package, individual item prices $\alpha\in\mathbb{R}_+^N$, and a radius of approximation $r \geq 1$. The bids can be used to impute a value function for non-empty packages for the core-determining engine according to the formula

$$v(S)=\begin{cases}\beta+\sum_{n\in S-T}\alpha_n-\sum_{n\in T-S}\alpha_n & \text{if } \max\left(|S-T|,|T-S|\right)\leq r\\ 0 & \text{otherwise}\end{cases}$$

where $|S-T|$ and $|T-S|$ are the numbers of elements in $S-T$ and $T-S$, respectively. Thus, the tuple (T,β,α,r) is understood to specify an offer of β for package *T* and adjustments for packages that are similar to *T*. Adding and/or deleting up to *r* items from the package *T* alters the bid by adding and subtracting the corresponding item prices. Adding and/or subtracting more than *r* items results in a zero bid (though it is should be evident from the logic that other specifications besides zero could also work here). The asset sale described above is a further simplification that restricts the sets *T* and the radius *r*. We denote by $\hat{V}$ the set of values that can be reported without any restrictions on *T* or *r*.

<u>Theorem 11</u>. Let $\Gamma_{\hat{V}}$ be the simplification of a standard core-selecting auction with reported valuations restricted to lie in $\hat{V}$. Then, regardless of the bidders' actual valuations, this mechanism has the best-reply closure property. (The same is true even when *r* is restricted, but not when *T* is restricted.)

The proof is a simple application of Theorem 3, because the simplification here is complete. To see why, fix some bidder *j* and strategies in $\hat{V}$ for the other bidders. Suppose there is some best reply report by *j* that wins some non-empty package *T* at price p_T. Let $\alpha_j \in \mathbb{R}_+^N$ be any vector with the property that for all $n \in T$, $\alpha_{jn} > \alpha_{j'n}$ for all other bidders $j' \neq j$ and for $n \notin T$, $\alpha_{jn} = 0$. Since the auction selects core allocations with respect to the reports, it chooses goods assignments to maximize total value. Therefore, the allocation selected by the original best-reply has a higher total value than any allocation that excludes *j*. So, *j* must still be a winner with the proposed bid. By construction, the value-maximizing outcome when *j* is included assigns package *T* to *j*. Since the mechanism is standard and *j*'s winning bid is unchanged, its price is also unchanged.

The preceding argument works for any value of *r*, so restrictions on *r* do not change the conclusion.

One interesting aspect of the affine approximation mechanisms is that they use something resembling prices to guide the allocation of items among the winning bidders. The idea of using item prices to guide package allocation has been repeatedly proposed in recent years. It is incorporated in the FCC's current package bidding algorithm and in the dynamic algorithms suggested by Porter, Rassenti, Roopnarine, and Smith (2003) and by Ausubel, Cramton, and Milgrom (2005). All of these mechanisms, however, impose upon prices the burden of guiding both the winner determination problem—which bidders should be in the winning set—and allocations of items among the winners.

The approximation mechanisms do not work that way: they attempt to utilize item prices to allocate goods among the winners but not by themselves to determine the set of winning bidders. The FCC's experiments with its package auction design shows that these item prices are highly unstable during the course of an ascending auction, increasing and decreasing by large amounts over time. In the perspective taken here, the proper item prices to guide the allocation of items among winners depends on the set of winners. If these are changing during an ascending package auction, then sharp swings in the supporting prices are to be expected.

The affine approximation mechanism with no restrictions on T or r may be useful in some settings with small number of items, but as the number of items grows large, they may admit too many coordination failure outcomes in which the number of packages implicitly bid by each bidder is too small. For some applications, one might require $r = N$, so that all bids are based on a single affine approximation of each bidder's value function. Such a mechanism makes computation easy and transparent and reduces size of the bid/report from something that is exponential in N to something that is linear in N. More generally, restricting T and/or requiring a wide radius of approximation r or using a better approximation than the affine one may be workable simplifications for some applications.

Small Bid Costs

The idea that bids costs are significant in package auctions even with relatively few items seems compelling—with $N = 10$ items, there are $2^N - 1 = 1023$ non-empty packages. Nevertheless, the best way to introduce these costs into the analysis is not obvious. One particularly simple alternative is to assume that costs are zero for simplifications that make the number of reports rise only linearly in N and the cost is otherwise prohibitive. By this standard, the affine approximation auctions described above are zero cost mechanisms, while full menu auctions are prohibitively costly. If the bid reductions are left to the bidders, there are many equilibria involving coordination failures, where packages in the efficient allocation receive no bid at all.

Another approach to bidding costs, more consistent with the treatment of sponsored search auctions above, is to assume that there is some small cost $c > 0$ of reporting each number. The difficulties this poses for equilibrium analysis are most simply illustrated by considering the case of a single item for sale: N=1. Suppose there are two bidders: a high value bidder 1 with value v_1 and a low value bidder 2 with value v_2. In the second-price auction in this case, the only full-information equilibrium has bidder 1 bid v_1 while bidder 2 bids zero, so the seller's revenue is zero. The first-price auction has no full information pure equilibrium when bid costs are small and positive. For if there were such an equilibrium and the equilibrium price were less than v_2, then both bidders 1 and 2 would enter, leading to a higher price than v_2. Alternatively, if the equilibrium price were v_2 or higher, then only bidder 1 would enter, so the price would be zero. It seems sensible for this case to model small bid costs by focusing on a pure price that is a limit of mixed strategy equilibria with random participation by bidder 2. This limiting price must be v_2, for if the bidder 2 randomizes about entry, its equilibrium profit must be zero, so the probability that a bid of $v_2 - c - k$ wins can be no more than c / k.

This analysis points to a revenue advantage to using first-price auctions rather than second-price auctions when bid costs are positive but small. Day and Milgrom (2007) reach an opposite conclusion using a different idea, namely, that it is cheaper to bid straightforwardly than to base each bid on a strategic calculation, so that the cost of bidding is less in a second-price auction. This may also encourage more entry. Neither of these effects appears in our full-information equilibrium analysis, but that is an outcome of the particular and extreme assumptions required for such an analysis. Our model is not well suited to assess the comparative importance of these competing effects, but it does succeed in highlighting a new and potentially significant effect.

V. Conclusion

That simplicity is desirable seems uncontroversial, yet there has been little discussion about what "simplicity" means, what advantages it conveys, or what kinds of problems result from inappropriate simplifications. Here, we tackle those questions by defining a *simplified mechanism* to be a direct mechanism but with a restriction on the set of permissible reports or bids. Packaging multiple goods into lots is a simplification in just this sense.

A common simplification involves conflating two or more distinct assignments and applying the same bid to both. Practically every real auction involves some conflation, for example because the bid typically depends only on the items acquired by the bidder and not on the assignment of the remaining goods. Sponsored search involves a further conflation because the same price per click must be offered for all ads for a particular search term regardless of the position of the ad on the search page (and, sometimes, to similar search terms as judged by an automated algorithm). Treasury bills with different serial numbers are such obvious candidates to be conflated that one might overlook that requiring the same bid for each bill with the same face value is a simplification. It would, of course, be possible to distinguish bills by serial number, but conflating bids to eliminate the distinction conveys the same advantages as the similar restriction in of sponsored search: it eliminates low revenue equilibria (including both pure and mixed equilibria).[10] Conflation is also used in certain electrical power auctions, when "zones" are established within which power or capacity is treated as a single undifferentiated commodity. This may be done even though substitution among power sources or sinks within a zone is imperfect.

One implication of all these examples is that conflation can increase competition for a set of goods by forcing a bid on one to be a bid on all. Yet not all conflations work equally well. In daily electrical power markets, the system operator typically acquires both base load generation capacity and load-following *regulation*—the latter is capacity that can produce power that follows the "load" (the power demanded) as it fluctuates from minute to minute. In California, losing bids to supply regulation were for a period not applied also as bids to supply base load capacity. In this case, a better simplification conflates asymmetrically: a bid for regulation should also count as a bid for base load capacity. The old system sometimes deprived the market of actually available base-load supply resulting in unnecessarily high prices.[11] This California case highlights both the tendency of practitioners to adopt simplified auction designs and the importance of choosing the right simplification.

---

[10] A T-bill auction that conflates bills with different serial numbers satisfies that best-reply closure property holds even in *mixed* strategies, so theorem 2 applies. To illustrate the advantage of conflations when bidding is costly, suppose there are N bills and $N+1$ bidders, that each bill is worth 1 to each bidder, and that each bidder can costlessly bid for one bill but incurs a cost to bid for two or more. Then, the unique Nash strategy equilibrium of the simplified first-price mechanism with a zero minimum bid has revenue of N, but no equilibrium of the auction for N individual items has revenue greater than 1.

[11] To illustrate how this can happen at equilibrium, imagine that demand fluctuates between 1 and 2 units and that there are three suppliers, each capable of supplying one unit and two capable of supplying regulation services by following the load fluctuations. If the two markets for base load and regulation are run separately and simultaneously, then there is a necessarily a single bidder in one of the markets.

In our theoretical account, simplification can have several advantages. First, in multi-product auctions, simplification can save costs by obviating the need to bid separately for all the possible alternatives. Second, in the same setting, simplification can improve performance because, if bidders can decide what reports to make, they may make too few bids or bid for the wrong packages, damaging efficiency and reducing revenues. For sponsored search auctions with positive bid costs and without simplification, we found that every full-information equilibrium entails zero seller revenues (for both the Vickrey design and the series of second-price auctions); in contrast, there are no zero-revenue equilibria in suitably simplified versions of these auctions. Third, even when bidding costs are zero, the full direct mechanism can have multiple Nash equilibria, some of which entail undesired outcomes. The Simplification Theorem applies to this zero-cost case, asserting that a simplification satisfying the best-reply closure property never introduces new equilibria, but may eliminate some equilibria by striking one of the equilibrium strategies.

Our theoretical account captures only some of the important aspects of simplified designs. It does not account for learning, which one might conjecture is faster and more precise in a simpler mechanism. It does not analyze the confusion that is created by complex mechanisms. It omits the resistance of bidders to participating in too complex a mechanism. Any of these features could be important.

Simplification is an essential aspect of practical mechanism design.

References

Ausubel, Lawrence, Peter Cramton, and Paul Milgrom. 2005. "The Clock-Proxy Auction: A Practical Combinatorial Auction Design," in *Combinatorial Auctions*. Peter Cramton, Yoav Shoham and Richard Steinberg eds. Cambridge, MA: MIT Press.

Ausubel, Lawrence and Paul Milgrom. 2002. "Ascending Auctions with Package Bidding." *Frontiers of Theoretical Economics*, 1:1, pp. Article 1.

Ausubel, Lawrence and Paul Milgrom. 2005. "The Lovely but Lonely Vickrey Auction," in *Combinatorial Auctions*. Peter Cramton, Yoav Shoham and Richard Steinberg eds. Cambridge, MA: MIT Press.

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

Cantillon, Estelle and Martin Pesendorfer. 2005. "Auctioning Bus Routes: The London Experience," in *Combinatorial Auctions*. Peter Cramton, Yoav Shoham and Richard Steinberg eds. Princeton: Princeton University Press.

Cramton, Peter, Yoav Shoham, and Richard Steinberg. 2005. *Combinatorial Auctions*. Cambridge, MA: MIT Press.

Day, Robert W. and Paul Milgrom. 2007. "Core-Selecting Package Auctions." *International Journal of Game Theory*, Forthcoming.

Edelman, Benjamin, Michael Ostrovsky, and Michael Schwartz. 2007. "Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords." *American Economic Review*, 97:1.

Koopmans, Tjalling and Martin Beckmann. 1957. "Assignment Problems and the Location of Economic Activities." *Econometrica*, 25:1, pp. 53-76.

Milgrom, Paul. 2004. *Putting Auction Theory to Work*. Cambridge: Cambridge University Press.

Porter, David, Stephen Rassenti, Anil Roopnarine, and Vernon Smith. 2003. "Combinatorial Auction Design." *Proceedings of the National Academy of Sciences,* 100, pp. 11153-57.

Rothkopf, Michael. 2007. "Thirteen Reasons Why the Vickrey-Clarke-Groves Mechanism is Not Practical." *Operations Research*, 55:2, pp. 191-97.

Simon, Leo K. and William R. Zame. 1990. "Discontinuous Games and Endogenous Sharing Rules." *Econometrica*, 58, pp. 861-72.

Varian, Hal R. 2006. "Position Auctions." Working paper.