

Core-Selecting Package Auctions¹

Robert Day and Paul Milgrom²

April 3, 2007

Auctions that select core allocations with respect to reported values generate competitive levels of sales revenues at equilibrium and limit bidder incentives to use shills. Among core-selecting auctions, the ones that minimize seller revenues also maximize incentives for truthful reporting, produce the Vickrey outcome when that lies in the core and, in contrast to the Vickrey auction, and create no incentive for a seller to exclude qualified bidders. Core-selecting auctions are related to and share properties with stable matching mechanisms.

Keywords: *core, stable matching, marriage problem, auctions, core-selecting auctions, menu auctions, proxy auctions, package bidding, combinatorial bidding, incentives, truncation strategies.*

JEL Codes: *D44, C78*

I. Introduction

Recent years have seen several new and important applications of matching procedures in practical applications, including school assignments in New York and Boston and new designs for life-saving organ exchanges. The mechanisms that have been adopted, and sometimes even the runner-up mechanisms, are stable matching mechanisms. Recall that *stable matches* are matches with the property that no individual can do better by staying unmatched and no pair can both do better by matching to one another. Since pairs are the only significant coalitions in this

¹ This paper evolved from Milgrom (2006), which reported a portion of Milgrom's Clarendon lectures for 2005. The authors subsequently discovered that versions of Theorem 1 appeared both in that paper and one produced independently by Day and Raghavan (2006). We have collaborated on this revision; in particular, nearly all of the material on shill bidding is new.

² Milgrom received financial support for this research from National Science Foundation under grant ITR-0427770. We thank Roger Myerson for suggesting the connection to Howard Raiffa's observations about bargaining, Yeon-Koo Che for comments on an earlier draft, and Manuj Garg for proofreading.

theory, stable matches are a kind of core allocation. *Stable matching mechanisms* are direct mechanisms that select a stable match whenever parties report truthfully; the definition does not require that the mechanism be incentive-compatible.

Evidence suggesting that stable matching mechanisms remain in use long after unstable mechanisms have been abandoned is found both in empirical studies (Roth and Xing (1994)) and in laboratory experiments (Kagel and Roth (2000)). If stable mechanisms actually lead to stable matches, then these mechanisms have the important practical advantage that no couple that would prefer to renege after the mechanism is run in favor of some alternative pairing, because no such agreement can be better for both members of the couple than the outcome of a stable matching mechanism. Even for a stable mechanism, with enough uncertainty, there might be pairs that could increase their expected payoffs by matching in advance, but the resulting unstable match would be vulnerable to defections by parties who might find a better alternative.

A similar analysis applies to core-selecting auction mechanisms. An individually rational outcome is in the core of an auction game if and only if there is no group of bidders who would strictly prefer an alternative deal that is also strictly better for seller. Consequently, an auction mechanism that delivers core allocations has the advantage that there is no individual or group that would want either to renege *after* the auction is run in favor of some allocation that is feasible for it and the any non-core agreement made before the auction risks being unwound afterwards.

For both matching and auction mechanisms, the preceding arguments have full force only if the procedures actually result in stable or core allocations, which in turn depends on the participants' strategies. Casual evidence suggests that participant behavior in real mechanisms varies widely from naïve to sophisticated, and the most sophisticated strategies involve not only reports in the mechanism but also decisions about whether to participate, to retrade after the mechanism has run, to renege on deals, to participate as a single player or as several, or even to attempt to alter the timing or rules of the mechanism itself.

It is customary in mechanism design theory to emphasize incentives for truthful reporting and to model incentive constraints on a par with resource constraints. However, because incentive constraints incorporate an imperfect behavioral theory and often conflict with other desiderata, it can be useful to replace the approach that gives primacy to incentive constraints by one that evaluates trade-offs among incentives and other desiderata. The conflicts between incentive-compatibility and other criteria are sometimes stark. It is known, for example, there exists no strategy-proof two-sided stable matching mechanism³ and that the only strategy-proof efficient auction mechanism, which is the (“generalized”) Vickrey auction,⁴ suffers from several severe practical drawbacks, some of which are described below.⁵ The task of those working in practical applications of mechanism design is to illuminate trade-offs by asking such questions as: what is the best we can do on criterion x for any given levels of performance on criteria y and z ?

The modern literature makes some attempts to account for multiple performance criteria even when incentives are less than perfect. Consider, for example, the basic two-sided matching problem, commonly called the *marriage problem*, in which men have preferences over women and women have preferences over men. The literature often treats stability of the outcome as the primary concern while still evaluating the incentive properties of the mechanism. In the marriage problem, there always exists a unique *man-optimal* match and a unique *woman-optimal* match.⁶ The man-optimal mechanism, which is the direct mechanism that always selects the man-optimal match, is strategy-proof for men but not for women⁷ and the reverse is true for the woman-

³ Roth (1982).

⁴ A result of Green and Laffont (1979), as extended by Holmstrom (1979), shows that for any path-connected set of valuations (for environments with quasi-linear preferences), the only strategy-proof direct auction mechanism that selects total-value-maximizing choices is the Vickrey mechanism

⁵ For a more thorough treatment, see Ausubel and Milgrom (2005).

⁶ As Gale and Shapley first showed, there is a stable match that is Pareto preferred by all men to any other stable match, which they called the “man optimal” match.

⁷ Hatfield and Milgrom (2005) identify the conditions under which strategy-proofness extends to cover the college admissions problem, in which one type of participant (“colleges”) can accept multiple applicants, but the other kind (“students”) can each be paired to only one college. Their analysis also covers problems in which wages and other contract terms are endogenous.

optimal mechanism. Properties such as these are typically reported as advantages of the mechanism,⁸ even though these incentives fall short of full strategy-proofness. Even when strategy-proofness fails, finding profitable deviations may be so hard that many participants find it most attractive just to report truthfully. A claim of this sort has been made for the pre-1998 algorithm used by National Resident Matching Program, which was not strategy-proof for doctors, but for which few doctors could gain at all by misreporting and for which tactical misreporting was fraught with risks (Roth and Peranson (1999)).⁹

The analysis of multiple criteria is particularly important for the design of *package auctions* (also called “combinatorial auctions”), which are auctions for multiple items in which bidders can bid directly for non-trivial subsets (“packages”) of the items being sold, rather than being restricted to submit bids on each item individually. In these auctions, several criteria besides incentive compatibility merit the attention of a practical mechanism designer. Revenues are an obvious one. Auctions are commonly run by an expert auctioneer on behalf of the actual seller and any failure to select a core allocation with respect to reported values implies that there is a group of bidders who have offered to pay more in total than the winning bidders, yet whose offer has been rejected. Imagine trying to explain such an outcome to the actual seller or, in a government sponsored auction, to a skeptical public!¹⁰ Monotonicity of revenues with respect to participation is another important property of auction mechanisms, because its failure could allow a seller to increase sales revenues by disqualifying bidders after the bids are received.¹¹ Another

⁸ For example, see Abdulkadiroglu, Pathak, Roth and Sonmez (2005).

⁹ There is quite a long tradition in economics of examining approximate incentives in markets, particularly when the number of participants is large. An early formal analysis is by Roberts and Postlewaite (1976).

¹⁰ McMillan (1994) describes how heads rolled when second-price auctions were used to sell spectrum rights in New Zealand and the highest bid was sometimes orders of magnitude larger than the second highest bid.

¹¹ It is quite common in auctions for final evaluations of bidder qualifications to be made after bids are received to ensure the winning bidder’s ability to close the deal. Sellers may carefully study financing and regulatory constraints before accepting a bid. In the US radio spectrum auctions, bidders typically submit a “short form application” before each auction and, after the bidding but before the results are finalized, the

important desideratum is that a bidder should not profit by entering and playing as multiple bidders, rather than as a single one.¹²

We illustrate these three desiderata and how they fail in the Vickrey auction with an example of two identical items are for sale. The first bidder wants both items and will pay up to 10 for the pair; it has zero value for acquiring a single item. The second and third bidders each have values of 10 for either one or two items, so their marginal values of a second item are zero. The Vickrey auction outcome assigns the items to the second and third bidders for prices of zero. Given that any of the three bidders would pay 10 for the pair of items, a zero price is surely too low: that is the low revenue problem.¹³ Also, the seller could increase its sales revenue by disqualifying bidder 3, thereby raising the total Vickrey price to 10. This illustrates the disqualification problem created by revenue non-monotonicity. Finally, if the second and third bidders are both controlled by the same player, then that player has reduced its price from 10 to zero by pretending to be two players rather than by bidding as a single entity: that is the shill bidding problem. These vulnerabilities are so severe that practical mechanism designers must investigate when and whether relaxing incentive compatibility can alleviate the problems.

We have discussed matching and package auction mechanisms together not only because they are two of the currently mostly active areas of practical mechanism design but also because there are some remarkable parallels between their equilibrium theories. One parallel connects the cases where the workers in the match are substitutes for hospital and when the goods in the

winning bidders make an additional cash deposit and submit a “long form” that is checked in detail to ensure their qualifications to buy.

¹² Yokoo, Sakurai and Matsubara (2004) were the first to emphasize the importance of “false name bidding” and how it could arise in the anonymous environment of Internet auctions. The problem they identified, however, is broader than just anonymous Internet auctions. For example, in the US radio spectrum auctions, several of the largest corporate bidders (including AT&T, Cingular, T-Mobile, Sprint, and Leap Wireless) have at times had contracts with or financial interests in multiple bidding entities in the same auction, enabling strategies that would not be possible for a single, unified bidder.

¹³ In this case, a core outcome requires that the seller receive at least 10. The bidder optimal outcomes—core outcomes that are not Pareto-dominated for bidders—are the outcomes in which 2 and 3 are the winning bidders, each pays a price between zero and ten, and the total payments are exactly ten.

auction are substitutes for the bidders. In these cases, the mechanism that selects the doctor-optimal match is *ex post* incentive-compatible for doctors and a mechanism, the ascending proxy auction of Ausubel and Milgrom (2002), which selects a bidder-optimal allocation (a core allocation that is Pareto optimal for bidders), is *ex post* incentive-compatible for bidders.¹⁴

A second important connection is the following one: for every stable match x and every stable matching mechanism, there exists an equilibrium in which each player adopts a certain *truncation strategy*, according to which it truthfully reports its ranking of all the outcomes at which it is not matched, but reports that it would prefer to be unmatched rather than to be assigned an outcome worse than x . What is remarkable about this theorem is that *one single profile of truncation strategies is a Nash equilibrium for every stable matching mechanism*. We will find that a similar property is true for core-selecting auctions, but with one difference. In matching mechanisms, it is usual to treat all the players as strategic, whereas in auctions it is not uncommon to treat the seller differently, with only a subset of the players—the *bidders*—treating as making decisions strategically. We are agnostic about whether to include the seller as a bidder or even whether to include all the buyers as strategic players. Regardless of how the set of strategic players is specified, we find that for every allocation on the Pareto-frontier of the core for the players who report strategically, there is a single profile of truncation strategies that is an equilibrium profile for every core-selecting auction.¹⁵

The preceding results hinge on another similarity between package auctions and matching mechanisms. In any stable matching mechanism or core-selecting auction and given any reports by the other players, a player's best reply achieves its maximum core payoff or best stable match given its actual preferences and the reported preferences of others. For auctions, there is an

¹⁴ This is also related to results on wage auctions in labor markets as studied by Kelso and Crawford (1982) and Hatfield and Milgrom (2005), although these models do not employ package bidding.

¹⁵ These truncation strategies also coincide with what Bernheim and Whinston (1986) call “truthful strategies” in their analysis of a “menu auction,” which is a kind of package auction.

additional interesting connection: the maximum core payoff is exactly the Vickrey auction payoff.

Next are the inter-related results about incentives for *groups* of participants. Given a core-selecting auction, the incentives for misreporting are minimal for individuals in a particular group S if and only if the mechanism selects an S -best core allocation. If there is a unique S -best allocation, then truthful reporting by members of coalition S is an *ex post* equilibrium. This is related to the famous result from matching theory (for which there always exists a unique man-optimal match and a unique woman-optimal match) that it is an *ex post* equilibrium for men to report truthfully in the man-optimal mechanism and for women to report truthfully in the woman-optimal mechanism.

Another result is that any auction that minimizes the seller's revenue among core allocations results in seller revenue being a non-decreasing function of the bids. As argued above, revenue-monotonicity of this sort is important because, without it, a seller might have an incentive to disqualify bids or bidders to increase its revenues and a bidder might have an incentive to sponsor a skill, whose bids reduce prices.

The remainder of this paper is organized as follows. Section II formulates the package auction problem. Section III characterizes core-selecting mechanisms in terms of revenues that are never less than Vickrey revenues, even when bidders can use skills. Section IV introduces definitions and notation and introduces the theorems about best replies and full information equilibrium. Section V states and proves the theorem about the core-selecting auctions with the smallest incentives to misreport. Section VI shows that the revenue-minimizing core-selecting auction is revenue-monotonic. Various corresponding results for the marriage problem are developed in section VII, while section VIII concludes.

II. Formulation

We denote the seller as player 0, the bidders as players $j = 1, \dots, J$, and the set of all players by N . Each bidder j has quasi-linear utility and a finite set of possible packages X_j . Its value associated with any feasible package $x_j \in X_j$ is $u_j(x_j) \geq 0$. For convenience, we formulate our discussion mainly in terms of bidding applications, but the same mathematics accommodates much more, including some social choice problems. In the central case of package bidding for predetermined items, x_j consists of a package of items that the bidder may buy. For procurement auctions, x_j could also usefully incorporate information about delivery dates, warranties, and various other product attributes or contract terms. Among the possible packages for each bidder is the null package, $\emptyset \in X_j$ and we normalize so that $u_j(\emptyset) = 0$.

For concreteness, we focus on the case where the auctioneer is a seller who has a feasible set $X_0 \subseteq X_1 \times \dots \times X_J$ with $(\emptyset, \dots, \emptyset) \in X_0$ —so the no sale package is feasible for the seller—and a valuation function $u_0 : X_0 \rightarrow \mathbb{R}$ normalized so that $u_0(\emptyset, \dots, \emptyset) = 0$. For example, if the seller must produce the goods to be sold, then u_0 may be the auctioneer-seller's variable cost function.

For any coalition S , a goods assignment \hat{x} is *feasible* for coalition S , written $\hat{x} \in F(S)$, if (1) $\hat{x} \in X_0$ and (2) for all j , if $j \notin S$ or $0 \notin S$, then $\hat{x}_j = \emptyset$. That is, a bidder can have a non-null assignment when coalition S forms only if that bidder and the seller are both in the coalition.

The *coalition value function* or *characteristic function* is defined by:

$$w_u(S) = \max_{x \in F(S)} \sum_{j \in S} u_j(x_j) \quad (1)$$

In a *direct auction mechanism* (f, P) , each bidder j reports a valuation function \hat{u}_j and the profile of reports is $\hat{u} = \{\hat{u}_j\}_{j=1}^J$. The outcome of the mechanism, $(f(\hat{u}), (P_j(\hat{u}))) \in (X_0, \mathbb{R}_+^J)$, specifies the choice of $x = f(\hat{u}) \in X_0$ and the payments $p_j = P_j(\hat{u}) \in \mathbb{R}_+$ made to the seller by

each bidder j . The associated payoffs are given by $\pi_0 = u_0(x) + \sum_{j \neq 0} p_j$ for the seller and $\pi_j = u_j(x) - p_j$ for each bidder j . The payoff profile is individually rational if $\pi \geq 0$.

A *cooperative game* (with transferable utility) is a pair (N, w) consisting of a set of players and a characteristic function. A payoff profile π is feasible if $\sum_{j \in N} \pi_j \leq w(N)$, and in that case it is associated with a feasible allocation. An *imputation* is a feasible, non-negative payoff profile. An imputation is in the *core* if it is efficient and unblocked:

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

A direct auction mechanism (f, P) is *core-selecting* if for every report profile \hat{u} , $\pi_{\hat{u}} \in Core(N, w_{\hat{u}})$. Since the outcome of a core-selecting mechanism must be efficient with respect to the reported preferences, we have the following:

Lemma 1. For every core-selecting mechanism (f, P) and every report profile \hat{u} ,

$$f(\hat{u}) \in \arg \max_{x \in X_0} \sum_{j \in N} \hat{u}_j(x_j) \quad (3)$$

The payoff of bidder j in a Vickrey auction is the bidder's marginal contribution to the coalition of the whole. In cooperative game notation, if the bidders' value profile is u , then bidder j 's payoff is $\bar{\pi}_j = w_u(N) - w_u(N - j)$.¹⁶

III. Revenues and Shills: Necessity of Core-Selecting Auctions

We have argued that the revenues from the Vickrey outcome are often too low to be acceptable to auctioneers. In order to avoid biasing the discussion too much, in this section we treat the Vickrey revenues as a just-acceptable lower bound and ask: what class of auctions have the properties that, for any set of reported values, they select the total-value maximizing outcome

¹⁶ A detailed derivation can be found in Milgrom (2004).

and lead always to revenues no lower than the Vickrey revenues, even when bidders may be using skills? Our answer will be: exactly the class of core-selecting auctions.

In standard fashion, we call any mechanism with the first property, namely, that the auction selects the total-value-maximizing outcome, “efficient.”

Theorem 1. An efficient direct auction mechanism has the property that no bidder can ever earn more than its Vickrey payoff by disaggregating and bidding with skills if and only if it is a core-selecting auction mechanism.

Proof. Fix a set of players (seller and bidders) N , let w be the coalitional value function implied by their reported values, and let π be the players’ vector of reported payoffs. Efficiency means $\sum_{j \in N} \pi_j = w(N)$. Let $S \subseteq N$ be a coalition that excludes the seller. These bidders could be skills. Our condition requires that they earn no more than if they were to submit their merged valuation in a Vickrey auction, in which case the merged entity would acquire the same items and enjoy a total payoff of equal to their marginal contribution to the coalition of the whole:

$w(N) - w(N - S)$. Our restriction is therefore $\sum_{j \in S} \pi_j \leq w(N) - w(N - S)$. In view of efficiency, this holds if and only if $\sum_{j \in N - S} \pi_j \geq w(N - S)$. Since S was an arbitrary coalition of bidders, we have that for every coalition $T = N - S$ that includes the seller, $\sum_{j \in T} \pi_j \geq w(T)$.

Since coalitions without the seller have value zero and can therefore never block, we have shown that there is no blocking coalition. Together with efficiency, this implies that $\pi \in \text{Core}(N, w)$. \square

IV. Truncation Reports and Equilibrium

In the marriage problem, a *truncation report* refers to a reported ranking by person j that preserves the person’s true ranking of possible partners, but which may falsely report that some partners are unacceptable. For an auction setting with transferable utility, a truncation report is similarly defined to correctly rank all pairs consisting of a non-null goods assignment and a

payment but which may falsely report that some of these are unacceptable. When valuations are quasi-linear, a reported valuation is a truncation report exactly when all reported values of non-null goods assignments are reduced by the same non-negative constant. We record that observation as a lemma.

Lemma 2. A report \hat{u}_j is a truncation report if and only if there exists some $\alpha \geq 0$ such that for all $x_j \in X_j$, $\hat{u}_j(x_j) = u_j(x_j) - \alpha$.

Proof. Suppose that \hat{u}_j is a truncation report. Let x_j and x'_j be two non-null packages and suppose that the reported value of x_j is $\hat{u}_j(x_j) = u_j(x_j) - \alpha$. Then, $(x_j, u_j(x_j) - \alpha)$ is reportedly indifferent to $(\emptyset, 0)$. Using the true preferences, $(x_j, u_j(x_j) - \alpha)$ is actually indifferent to $(x'_j, u_j(x'_j) - \alpha)$ and so must be reportedly indifferent as well: $\hat{u}_j(x_j) - u_j(x_j) - \alpha = \hat{u}_j(x'_j) - u_j(x'_j) - \alpha$. It follows that $u_j(x'_j) - \hat{u}_j(x'_j) = u_j(x_j) - \hat{u}_j(x_j) = \alpha$.

Conversely, suppose that there exists some $\alpha \geq 0$ such that for all $x_j \in X_j$, $\hat{u}_j(x_j) \equiv u_j(x_j) - \alpha$. Then for any two non-null packages, the reported ranking of (x_j, p) is higher than that of (x'_j, p') if and only if $\hat{u}_j(x_j) - p \geq \hat{u}_j(x'_j) - p'$ which holds if and only if $u_j(x_j) - p \geq u_j(x'_j) - p'$. \square

We refer to the truncation report in which the reported value of all non-null outcomes is $\hat{u}_j(x_j) = u_j(x_j) - \alpha_j$ as the “ α_j truncation of u_j .”

In full information auction analyses since that of Bertrand (1883), auction mechanisms have often been incompletely described by the payment rule and the rule that the unique highest bid, when that exists, determines the winner. Ties often occur at Nash equilibrium, however, and the way ties are broken is traditionally chosen in a way that depends on bidders' values and not just on their bids. For example, in a first-price auction with two bidders, both bidders make the

same equilibrium bid, which is equal to the lower bidder's value. The analysis assumes that the bidder with the higher value is favored, that is, chosen to be the winner in the event of a tie. If the high value bidder were not favored, then it would have no best reply. As Simon and Zame (1990) have explained, although breaking ties using value information prevents this from being a feasible mechanism, the practice of using this tie-breaking rule for analytical purposes is an innocent one, because, for any $\varepsilon > 0$, the selected outcome lies within ε of the equilibrium outcome of any related auction game in which the allowed bids are restricted to lie on a sufficiently fine discrete grid.¹⁷

In view of lemma 1, for almost all reports, assignments of goods differ among core-selecting auctions only when there is a tie; otherwise, the auction is described entirely by its payment rule. We henceforth denote the payment rule of an auction by $P(\hat{u}, x)$, to make explicit the idea that the payment may depend on the goods assignment in case of ties. For example, a first-price auction with only one good for sale is any mechanism which specifies that the winner is a bidder who has made the highest bid and the price is equal to that bid. The mechanism can have any tie-breaking rule to be used so long as (3) is satisfied. In traditional parlance, the payment rule P defines an *auction*, which comprises a set of mechanisms.

Definition. \hat{u} is an equilibrium of the auction P if there is some core selecting mechanism (f, P) such that \hat{u} is a Nash equilibrium of the mechanism.

For any auction, consider a tie-breaking rule in which bidder j is *favored*. This means that in the event that there are multiple goods assignments that maximize total reported value, if there is one at which bidder j is a winner, then the rule selects such a one. When a bidder is favored, that bidder always has some best reply.

¹⁷ See also Reny (1999).

Theorem 2. Suppose that (f, P) is a core-selecting direct auction mechanism and bidder j is favored. Let \hat{u}_{-j} be any profile of reports of bidders other than j . Denote j 's actual value by u_j and let $\bar{\pi}_j = w_{\hat{u}_{-j}, u_j}(N) - w_{\hat{u}_{-j}, u_j}(N - j)$ be j 's corresponding Vickrey payoff. Then, the $\bar{\pi}_j$ truncation of u_j is among bidder j 's best replies in the mechanism and earns a payoff for j of $\bar{\pi}_j$. Moreover, this remains a best reply even in the expanded strategy space in which bidder j is free to use skills.¹⁸

Proof. Suppose j reports the $\bar{\pi}_j$ truncation of u_j . Since the mechanism is core-selecting, it selects individually rational allocations with respect to reported values. Therefore, if bidder j is a winner, its payoff is at least zero with respect to the reported values and hence at least $\bar{\pi}_j$ with respect to its true values.

Suppose that some report \hat{u}_j results in an allocation \hat{x} and a payoff for j strictly exceeding $\bar{\pi}_j$. Then, the total payoff to the other bidders is less than $w_{\hat{u}_{-j}, u_j}(N) - \bar{\pi}_j \leq w_{\hat{u}_{-j}, u_j}(N - j)$, so $N - j$ is a blocking coalition for \hat{x} , contradicting the core-selection property. This argument applies also when bidder j uses skills. Hence, there is no report yielding a profit higher than $\bar{\pi}_j$, even on the extended strategy space that incorporates skills.

Since reporting the $\bar{\pi}_j$ truncation of u_j results in a zero payoff for j if it loses and non-negative payoff otherwise, it is always a best reply when $\bar{\pi}_j = 0$.

Next, we show that the truncation report always wins for j , therefore yielding a profit of at least $\bar{\pi}_j$ so that it is a best reply. Regardless of j 's reported valuation, the total reported payoff to any coalition excluding j is at most $w_{\hat{u}_{-j}, \hat{u}_j}(N - j) = \max_{x=(\emptyset, x_{-j}) \in X_0} \sum_{i \in N-j} \hat{u}_i(x)$. If j reports the

¹⁸ Versions of this result were derived and reported independently by Day and Raghavan (2006) and by Milgrom (2006).

$\bar{\pi}_j$ truncation of u_j , then the maximum value is at least $\max_{x \in X_0} \left(\sum_{i \in N-j} \hat{u}_i(x) + u_j(x) \right) - \bar{\pi}_j = w_{\hat{u}_{-j}, u_j}(N) - \bar{\pi}_j$, which is equal to the previous sum by the definition of $\bar{\pi}_j$. Applying lemma 1 and the hypothesis that j is favored establishes that j is a winner. \square

Definition. An imputation π is *bidder optimal* if $\pi \in \text{Core}(N, u)$ and there is no $\hat{\pi} \in \text{Core}(N, u)$ such that for every bidder j , $\pi_j \leq \hat{\pi}_j$ with strict inequality for at least one bidder. (By extension, a feasible allocation is *bidder optimal* if the corresponding imputation is so.)

Next is one of the main theorems, which establishes a kind of equilibrium equivalence among the various core-selecting auctions. We emphasize, however, that the strategies require each bidder j to know the equilibrium payoff π_j , so what is being described is a full information equilibrium but not an equilibrium in the model where each bidder's own valuation is private information.

Theorem 3. For every valuation profile u and corresponding bidder optimal imputation π , the profile of π_j truncations of u_j is a full information equilibrium profile of every core selecting auction. The equilibrium goods assignment x^* maximizes the true total value $\sum_{i \in N} u_i(x_i)$, and the equilibrium payoff vector is π (including π_0 for the seller).

Proof. For any given core-selecting auction, we study the equilibrium of the corresponding mechanism that, whenever possible, breaks ties in (3) in favor of the goods assignment that maximizes the total value according to valuations u . If there are many such goods assignments, any particular one can be fixed for the argument that follows.

First, we show that no goods assignment leads to a reported total value exceeding π_0 . Indeed, let S be the smallest coalition for which the maximum total reported value exceeds π_0 .

By construction, the bidders in S must all be winners at the maximizing assignment, so

$$\pi_0 < \max_{x \in X_0, x_{-j} = \emptyset} u_0(x_0) + \sum_{i \in S-0} (u_i(x_i) - \pi_i) \leq w_u(S) - \sum_{i \in S-0} \pi_i .$$
 This contradicts

$\pi \in \text{Core}(N, w_u)$, so the winning assignment has a reported value of at most π_0 : $w_{\hat{u}}(N) \leq \pi_0$. If j

instead reports truthfully, it can increase the value of any goods allocation by at most π_j , so

$$w_{u_j, \hat{u}_{-j}}(N) \leq \pi_0 + \pi_j .$$

Next, we show that for any bidder j , there is some coalition excluding j for which the maximum reported value is at least π_0 . Since π is bidder optimal, for any $\varepsilon > 0$,

$(\pi_0 - \varepsilon, \pi_j + \varepsilon, \pi_{-j}) \notin \text{Core}(N, w_u)$. So, there exists some coalition S_ε to block it:

$$\sum_{i \in S_\varepsilon} \pi_i - \varepsilon < w_u(S_\varepsilon) .$$
 By inspection, this coalition includes the seller but not bidder j . Since this

is true for every ε and there are only finitely many coalitions, there is some S such that

$$\sum_{i \in S} \pi_i \leq w_u(S) .$$
 The reverse inequality is also implied because $\pi \in \text{Core}(N, w_u)$, so

$$\sum_{i \in S} \pi_i = w_u(S) .$$

$$\text{For the specified reports, } w_{\hat{u}}(S) = \max_{x \in X_0} \sum_{i \in S} \hat{u}_i(x_i) \geq$$

$$\max_{x \in X_0} u_0(x_0) + \sum_{i \in S-0} (u_i(x_i) - \pi_i) \geq w_u(S) - \sum_{i \in S-0} \pi_i = \pi_0 .$$
 Since the coalition value cannot

decrease as the coalition expands, $w_{\hat{u}}(N-j) \geq \pi_0$. By definition of the coalition value functions,

$$w_{\hat{u}}(N-j) = w_{u_j, \hat{u}_{-j}}(N-j) .$$

Using theorem 2, j 's maximum payoff if it responds optimally and is favored is

$$w_{u_j, \hat{u}_{-j}}(N) - w_{u_j, \hat{u}_{-j}}(N-j) \leq (\pi_0 + \pi_j) - \pi_0 = \pi_j .$$
 So, to prove that the specified report profile is an

equilibrium, it suffices to show that each player j earns π_j when these reports are made.

The reported value of the true efficient goods assignment is at least

$$\max_{x \in X_0} u_0(x_0) + \sum_{i \in N-0} (u_i(x_i) - \pi_i) = w(N) - \sum_{i \in N-0} \pi_i = \pi_0 .$$
 So, with the specified tie-breaking

rule, if the bidders make the specified truncation reports, the selected goods assignment will maximize the true total value.

Since the auction is core-selecting, each bidder j must have a reported profit of at least zero and hence a true profit of at least π_j , but we have already seen that these are also upper bounds on the payoff. Therefore, the reports form an equilibrium; each bidder j 's equilibrium payoff is precisely π_j , and that the seller's equilibrium payoff is $w_{\hat{a}}(N) - \sum_{i \in N-0} \pi_i = \pi_0$. \square

V. Minimizing Incentives to Misreport

Despite the similarities among the core-selecting mechanisms emphasized in the previous section, there are important differences among the mechanisms in terms of incentives to report valuations truthfully. For example, when there is only a single good for sale, both the first-price and second-price auctions are core selecting mechanisms, but only the latter is strategy-proof.

To evaluate various bidders' incentives to deviate from truthful reporting, we introduce the following definition.

Definition. The *incentive profile* for a core-selecting auction P at u is

$$\varepsilon^P = \{\varepsilon_j^P(u)\}_{j \in N-0} \text{ where } \varepsilon_j^P(u) \equiv \sup_{\hat{u}_j} u_j(f_j(u_{-j}, \hat{u}_j)) - P(u_{-j}, \hat{u}_j, f_j(u_{-j}, \hat{u}_j)) \text{ is } j\text{'s}$$

maximum gain from deviating from truthful reporting when j is favored.

Our idea is to minimize these incentives to deviate from truthful reporting, subject to selecting a core allocation. Since the incentives are represented by a vector, we use a Pareto-like criterion.

Definitions. A core-selecting auction P provides *suboptimal incentives* at u if there is some core selecting auction \hat{P} such that for every bidder j , $\varepsilon_j^{\hat{P}}(u) \leq \varepsilon_j^P(u)$ with strict inequality for some bidder. A core selecting auction *provides optimal incentives* if there is no u at which it provides suboptimal incentives.

Theorem 4. A core-selecting auction provides optimal incentives if and only if for every u it chooses a bidder optimal allocation.

Proof. Let P be a core-selecting auction, u a value profile, and π the corresponding auction payoff vector. From theorem 2, the maximum payoff to j upon a deviation is $\bar{\pi}_j$, so the maximum gain to deviation is $\bar{\pi}_j - \pi_j$. So, the auction is suboptimal exactly when there is another core-selecting auction with higher payoffs for all bidders, contradicting the assumption that π is bidder optimal. \square

Recall that when the Vickrey outcome is a core allocation, it is the unique bidder optimal allocation. So, Theorem 4 implies that any core selecting auction that provides optimal incentives selects the Vickrey outcome with respect to the reported preferences whenever that outcome is in the core for those reports. Moreover, because truthful reporting then provides the bidders with their Vickrey payoffs, theorem 2 implies the following.

Corollary. When the Vickrey outcome is a core allocation, then truthful reporting is an *ex post* equilibrium for any mechanism that always selects bidder optimal core allocations.

We note in passing that any incentive profile that can be achieved by any mechanism is replicated by the corresponding direct mechanism. There is a “revelation principle” for approximate incentives, so one cannot do better than the results reported in theorem 4 by looking over a larger class of mechanisms, including ones that are not direct.

VI. Monotonicity of Revenues

The core allocations with respect to the reports that minimize the seller’s payoff are all bidder optimal allocations, so a mechanism that selects those satisfies the conditions of theorem 4. That mechanism has another advantage as well: its revenues are non-decreasing in the bids.

Theorem 5. The seller's minimum payoff in the core with bidder values \hat{u} is non-decreasing in \hat{u} .

Proof. The seller's minimum payoff is:

$$\min_{\pi \geq 0} w_{\hat{u}}(N) - \sum_{i \in N-0} \pi_i \text{ subject to } \sum_{i \in S} \pi_i \geq w_{\hat{u}}(S) \text{ for all } S \subseteq N \quad (4)$$

The objective is an expression for π_0 ; it incorporates the equation $w_{\hat{u}}(N) = \sum_{i \in N} \pi_i$ which therefore can be omitted from the constraint set. The objective is increasing in $w_{\hat{u}}(N)$ and the constraint set shrinks as $w_{\hat{u}}(S)$ increases for any coalition $S \neq N$. Hence, the minimum value is non-decreasing in the vector $(w_{\hat{u}}(S))_{S \subseteq N}$. It is obvious that the coalitional values $w_{\hat{u}}(S)$ are non-decreasing in the reported values \hat{u} , so the result follows. \square

VII. Connections to the Marriage Problem

Even though Theorems 2-5 in this paper are proved using transferable utility, they all have analogs in the non-transferable utility marriage problem.

Consider Theorem 2. Roth and Peranson (1999) have shown for a particular algorithm in the marriage problem that any fully informed player can guarantee its best stable match by a suitable truncation report. That report states that all mates less preferred than its best achievable mate are unacceptable. The proof in the original paper makes it clear that their result extends to any stable matching mechanism, that is, any mechanism that always selects a stable match.

Here, in correspondence to stable matching mechanisms, we study core-selecting auctions. For the auction problem, Ausubel and Milgrom (2002) showed that the best payoff for any bidder at any core allocation is its Vickrey payoff. So, the Vickrey payoff corresponds to the best mate assigned at any stable match. Thus, the auction and matching procedures are connected not just by the use of truncation strategies as best replies but by the point of the truncation, which is at the player's best core or stable outcome.

Theorem 3 concerns Nash equilibrium. Again, the known results of matching theory are similar. Suppose the participants in the match in some set S^C play non-strategically, like the seller in the auction model, while the participants in the complementary set S , whom we shall call bidders, play Nash equilibrium. Then, for bidder-optimal stable match,¹⁹ the profile at which each player in S reports that inferior matches are unacceptable is a full-information Nash equilibrium profile of *every* stable matching mechanism and it leads to that S -optimal stable match. This result is usually stated using only men or women as the set S , but extending to other sets of bidders using the notion of bidder optimality is entirely straightforward.

For Theorem 4, suppose again that some players are non-strategic and that only the players in S report strategically. Then, if the stable matching mechanism selects an S -optimal stable match, then there is no other stable matching mechanism that weakly improves the incentives of all players to report truthfully, with strict improvement for some. Again, this is usually stated only for the case where S is the set of men or the set of women, and the extension does require introducing the notion of a bidder optimal match.

Finally, the last result states that increasing bids or, by extension, introducing new bidders increases the seller's revenue if the seller pessimal allocation is selected. The matching analog is that adding men improves the utility of each woman if the woman-pessimal, man-optimal match is selected—a result that is reported by Roth and Sotomayor (1990).

VIII. Conclusion

We motivated our study of core-selecting auctions by comparing them to stable matching mechanisms, which have been in long use in practice. Both in collected case studies and in the Kagel-Roth laboratory experiments, participants stopped using unstable matching mechanisms, preferring to make the best match they could by individual negotiations, even when congestion made that process highly imperfect. In contrast, participants continued to participate in stable

¹⁹ This is defined analogously to the bidder optimal allocation.

matching mechanisms for much longer, and many such mechanisms continue in use today. They have the practical advantage of being able to find stable allocations for the reported preferences, which would be a difficult task for them in the limited time typically available both in the experiments and in real markets.

These observations, however, compare only some particular stable and unstable matching mechanisms, so even the generalization to all matching mechanisms remains untested. If one can imagine experiments with other matching mechanisms, is there reason to not to extend that to variations in which the parties decide whether to bargain informally or to participate in an organized auction? Might such an auction experiment have a similar outcome?

Surely, whether the mechanisms being tested are matching or auction mechanisms, many details of both the design and the environment would matter. If the Vickrey auction were tested in experimental environments where the Vickrey payoffs lie is far outside the core, possibly where the revenue is always zero, it is hardly likely that subject sellers would voluntarily continue to use the auction. If, in analogy to the Kagel-Roth matching experiments, the seller had an opportunity to entertain direct offers instead, we would predict that the seller learn eventually to do just that. If the seller could choose among a set of auction designs, its preferred design might also depend on the prominence of shill bidder strategies. If, in analogy to the congestion of the experimental matching markets, the time allowed for bargaining makes it hard to find even a nearly efficient allocation, then there would be good reason for participants to participate voluntarily in the core-selecting auction, which selects such an allocation. Thus, both efficiency considerations and distributional considerations are likely to be important to the experimental outcome.

In auctions with N items for sale, the number of non-empty packages for which a bidder is called to report values is $2^N - 1$. That is unrealistically large for most applications if N is even a two-digit number. For the general case, Segal (2003) has shown that communications cannot be much reduced without severely limiting the efficiency of the result.

Although communication complexity is an important practical issue, it appears to differ in a qualitative way from the issues studied in this paper. In many environments, there is information about the kinds of packages that make sense and those can be incorporated as restrictions on bidding in an auction design. For example, an auctioneer may know that relaxing delivery constraints cannot increase costs, or that airport landing rights between 2:00-2:15 are valued similarly to ones between 2:15-2:30, or that complementarities in electrical generating result from costs saved by operating continuously in time, minimizing time lost when the plant is ramped up or down. Practical designs that take advantage of this structure can still be core-selecting mechanisms, where feasible allocations are subjected to the predetermined constraints.

If the problem of communication complexity can be solved, then core-selecting auctions appear to provide a practical alternative to the Vickrey design. The class includes the pay-as-bid “menu auction” design studied by Bernheim and Whinston (1986) as well as the ascending proxy auction studied by Ausubel and Milgrom (2002) and Parkes and Ungar (2000). Within this class, the auctions that select bidder-optimal allocations conserve as far as possible the advantages of the Vickrey design—matching the Vickrey auction’s *ex post* equilibrium property when there is a single good, or goods are substitutes, or most generally when the Vickrey outcome happens to lie in the core—and avoiding the low revenue and monotonicity problems of the Vickrey mechanism.

From the perspective of theory, the most interesting part of this analysis is that all of the main results about core-selecting auctions have analogues in the theory of stable matching mechanisms. The deep reasons for this similarity remain to be fully explored.

References

- Abdulkadiroglu, Atila, Parag Pathak, Alvin Roth and Tayfun Sonmez (2005). "The Boston Public School Match." *AEA Papers and Proceedings*: 368-371.
- Ausubel, Lawrence and Paul Milgrom (2002). "Ascending Auctions with Package Bidding." *Frontiers of Theoretical Economics* **1**(1): Article 1.
- Ausubel, Lawrence and Paul Milgrom (2005). "The Lovely but Lonely Vickrey Auction". *Combinatorial Auctions*. P. Cramton, Y. Shoham and R. Steinberg. Cambridge, MA, MIT Press.
- Bernheim, B. Douglas and Michael Whinston (1986). "Menu Auctions, Resource Allocation and Economic Influence." *Quarterly Journal of Economics* **101**: 1-31.
- Bertrand, J. (1883). "Théorie Mathématique de la Richesse Sociale." *Journal des Savants* **69**: 499-508.
- Day, Robert W. and S. Raghavan (2006). "Fair Payments for Efficient Allocations in Public Sector Combinatorial Auctions." *Management Science* **Forthcoming**.
- Green, Jerry and Jean-Jacques Laffont (1979). *Incentives in Public Decision Making*. North Holland: Amsterdam.
- Hatfield, John and Paul Milgrom (2005). "Matching with Contracts." *American Economic Review* **95**(4): 913-935.
- Holmstrom, Bengt (1979). "Groves Schemes on Restricted Domains." *Econometrica* **47**: 1137-1144.
- Kagel, John and Alvin Roth (2000). "The dynamics of reorganization in matching markets: A laboratory experiment motivated by a natural experiment." *Quarterly Journal of Economics*: 201-235.

- Kelso, Alexander and Vincent Crawford (1982). "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica* **50**: 1483-1504.
- McMillan, John (1994). "Selling Spectrum Rights." *Journal of Economics Perspectives* **8**: 145-162.
- Milgrom, Paul (2004). *Putting Auction Theory to Work*. Cambridge, Cambridge University Press.
- Milgrom, Paul (2006). "Incentives in Core-Selecting Auctions." Stanford University.
- Parkes, David and Lyle Ungar (2000). "Iterative Combinatorial Auctions: Theory and Practice." *Proceedings of the 17th National Conference on Artificial Intelligence*: 74-81.
- Reny, Philip (1999). "On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games." *Econometrica* **67**(5): 1029-1056.
- Roberts, John and Andrew Postlewaite (1976). "The Incentives for Price-Taking Behavior in Large Exchange Economies." *Econometrica* **44**(1): 115-129.
- Roth, Alvin E. (1982). "The Economics of Matching: Stability and Incentives." *Mathematics of Operations Research* **7**: 617-628.
- Roth, Alvin E. and Elliott Peranson (1999). "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review* **89**: 748-780.
- Roth, Alvin E. and Marilda Sotomayor (1990). *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge, Cambridge University Press.

- Roth, Alvin E. and X. Xing (1994). "Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions." *American Economic Review* **84**: 992-1044.
- Segal, Ilya (2003). "The Communication Requirements of Combinatorial Auctions". *Combinatorial Auctions*. P. Cramton, Y. Shoham and R. Steinberg, Princeton University Press.
- Simon, Leo K. and William R. Zame (1990). "Discontinuous Games and Endogenous Sharing Rules." *Econometrica* **58**: 861-872.
- Yokoo, Makoto, Yuko Sakurai and Shigeo Matsubara (2004). "The effect of false-name bids in combinatorial auctions: new fraud in internet auctions." *Games and Economic Behavior* **46**(1): 174-188.