Assignment Auctions

Paul Milgrom[1]
April 10, 2008

Abstract. We propose a new auction for selling goods that are close substitutes. The design is a direct mechanism that is compact, easy to implement, respects integer constraints, and is a "tight" simplification of a standard direct competitive mechanism. Connections are established between the assignment auction and the Vickrey auction, the uniform price auction for a single product, and an ascending multi-product clock auction.

Keywords: auctions, matching, substitutes, assignment, mechanism design

JEL Categories: D44, C78

I. Introduction

This paper describes and analyzes the *assignment auction* – a simplified mechanism for use in assigning and pricing multiple discrete varieties of a good. The analysis emphasizes simplification and, because this emphasis is new for mechanism design theory, we necessarily employ principles that are mostly untested and potentially controversial. What should not be controversial is that simplification is at the core of much of practical market design. In many real applications, the direct mechanisms studied in much of economic theory are far too complex to be useful.

[1] The assignment auction is a patent-pending invention of the author. Support for this research into the auction's properties was provided by National Science Foundation grant SES-0648293. Eduardo Perez and Clayton Featherstone made helpful suggestions on an earlier draft of this paper. Any opinions expressed here are those of the author alone.

The problem of complexity is endemic to both matching and auction mechanisms. For example, in the National Resident Matching Program (NRMP) which places doctors into hospital residency programs (Roth and Peranson (1999)), a hospital that interviews fifty candidates in the hopes of employing ten could be required by a full direct mechanism to rank all subsets of ten or less, resulting in a rank-order list of length approximately 1.3×10^{10}. The recently completed FCC auction 73 sold 1090 radio spectrum licenses, and a value report for all non-empty collections of those is a vector of dimension $2^{1090}-1$. Direct mechanisms that demand a full reporting of preferences are useless in these sorts of settings.

In practice, both matching and auction mechanisms employ simplifications to reduce the burden on participants. Hospitals in the NRMP report simple rank order lists of individual candidates, rather than sets of candidates, and the NRMP algorithm utilizes that information by employing a simplifying assumption about preferences – that they are "responsive." Responsive preferences exemplify one important heuristic principle – that a simplification should permit a good approximation of actual preferences in a large set of cases.

Ascending or descending auctions have been employed or proposed for many commercial applications in which similar but distinct goods are being sold. The goods could be radio spectrum licenses to use different but nearby frequencies, landing rights in nearby time slots at a busy airport, or commodities available at different locations or times, with different amounts of processing, or subject to different delivery guarantees. Assuming that goods are substitutes, simultaneous ascending or descending auctions of various sorts are potentially well suited to finding efficient or stable allocations or competitive prices (Kelso and Crawford (1982), Gul and Stacchetti (2000), Milgrom (2000), Ausubel (2004)), but implementing the direct mechanism which corresponds to these dynamical auctions would require a compact way to express all substitutable preferences. No compact representation of that set of preferences is currently available.

Two special characteristics of many of the practical applications suggest the possibility of a new, simplified auction design. First, it is commonly true that when different versions of a good are substitutable at all for a particular bidder, the rate of substitution is one-for-one. For example, a cement purchaser may wish to buy some quantity of cement and may be prepared to pay more to a supplier located closer to the point of use, but the number of tons needed may still be fixed independently of the source: substitution is one-for-one. A northern California electric utility may purchase power at the Oregon border or from southern California, subject to transmission constraints on each. Or, a cereal maker may be able to substitute bushels of grain today for bushels tomorrow by storing the grain in a suitable facility. These are all examples of auctions-to-buy, but a similar structure is also common in auctions-to-sell, as when a manufacturer can deliver several versions of the same processed good. In each case, substitution possibilities are typically limited, but when substitution is possible at all, it involves at least approximately one-for-one substitution among various versions of a good.[2]

Second, buyers and sellers often find it helpful or necessary to respect integer constraints, because many commodities are most efficiently shipped by the truckload or

[2] In hospital residency programs with a limited number of slots, the substitution among doctors is also one-for-one. The assignment auction may thus also be suitable for that application, provided that wages are made endogenous. An alternative auction for the medical match has been proposed by Crawford (2004).

container-load. A practical resource allocation mechanism may need to respect such integer constraints.

Often, an important objective of formal market mechanisms is to identify market-clearing prices. Sometimes, there are regulatory reasons for this objective. Other times, there may be organizational reasons. For example, if the seller is an agricultural supply cooperative, the auction price may be used both to apportion revenues fairly among sellers and to determine other transfers in the organization. When finding market clearing prices is a main goal, the mechanism can be conceived as a simplification of an abstract market mechanism in which each bidder reports a closed, convex-valued demand correspondence and the auctioneer selects an equilibrium consisting of a price vector and allocation.

To ease comparisons below, it is convenient to focus on a competitive mechanisms with the property that, when there are multiple equilibrium prices, the selected equilibrium price vector p is *minimal*, that is, there is no other equilibrium price vector p' satisfying $p' < p$. Given the selected price p, we will assume throughout that if there are multiple equilibrium allocations consistent with price p, then the mechanism selects among them using some fixed criterion F, for example, minimizing some strictly convex function on the convex, compact set of such allocations. We denote the resulting fully specified mechanism by μ_F^{CE}.

Simplification restricts the set of preferences or demands that a bidder can report in the auction.[3] This creates a potential problem, because profiles of strategies that were

[3] The formulation is already a simplification because it assumes preferences and restricts them in further ways. That simplification, which sets the context for our analysis, is not itself analyzed here.

not equilibria of the original game can become equilibria in some simplified games. Milgrom (2007) proposes the principle that a simplification should be *tight*, meaning that regardless of how participant preferences are specified within a certain relevant set, there should be no pure strategy Nash equilibrium of the simplification that is not also a Nash equilibrium of the original mechanism. To formalize tightness and related notions, we need some notation and definitions.

II. Preliminaries

We study environments with a set of possible *outcomes* X. A *mechanism* $\mu=(N,S,\omega)$ is a triple consisting of a set of players, a set of strategies profiles, ($S=S_1\times\ldots\times S_N$), and an outcome function $\omega:S\to X$. Adding a preference profile $(\succ_n)_{n\in N}$ for outcomes, one obtains a game with perfect information, which we denote by $\left(\mu,(\succ_n)_{n\in N}\right)$. In our application, outcomes will involve goods assignments and payments. We assume that each participant's preferences are represented by a utility function that depends only on its own assignment and payment and is decreasing and continuous in the amount of the payment; we call these *admissible* preferences.

<u>Definitions</u>.

1. The mechanism $\mu'=(N,S',\omega')$ is a *simplification* of $\mu=(N,S,\omega)$ (and μ is an *extension* of μ') if $S'\subseteq S$ and $\omega'=\omega_{|S'}$.

2. The simplification $\mu'=(N,S',\omega')$ of $\mu=(N,S,\omega)$ is *tight* if for all admissible preference profiles $(\succ_n)_{n\in N}$, every pure-strategy Nash

equilibrium profile of the simplified game $\left(\mu',(\succ_n)_{n\in N}\right)$ is also a Nash equilibrium of the extended game $\left(\mu,(\succ_n)_{n\in N}\right)$.

3. The mechanism $\mu'=(N,S',\omega')$ has the best-reply closure property with respect to its extension $\mu=(N,S,\omega)$ if $\left(\forall n\in N\right)\left(\forall s'_{-n}\in S'_{-n}\right)\left(\forall(\succ_m)_{m\in N}\right)$ $\left(\exists s'_n\in S'_n\right)\ s'_n\in \arg\max_{s_n\in S_n}\succ_n \omega(s_n,s'_{-n})$.

In words, these definitions say that a *simplification* is a mechanism with a more restricted strategy set than the original mechanism; it is *tight* if the restriction introduces no new pure strategy equilibria; and, it satisfies the *best-reply closure property* if for any profile of play by *j*'s competitors that entails only simplified strategies, *n* has a best reply in S_n that is also an element of S'_n.

Theorem 1. If the simplification μ' of μ has the best-reply closure property, then the simplification is tight.

Milgrom (2007) gives the first statement and proof of this theorem. I argue there that Google's generalized second price auction (Edelman, Ostrovsky, and Schwartz (2007), Varian (2006)) is a simplified auction in which bidders are allowed to name only a single price, rather than to bid a separate price for each ad position on the page. Incorporating the bidders' cost of making bids by refining equilibrium to require that bidders select only best replies with the fewest positive bids, the *extension* of Google's position auction has zero revenues in every refined equilibrium. In contrast, the actual, simplified mechanism ensures that active bidders make positive bids for all positions and guarantees positive revenue in every equilibrium.

III. Simple Assignment Auctions

In the assignment auction, we suppose that there are K varieties of a good which are offered for sale in positive quantities $(q_1,...,q_K)$. Bidders submit a finite set of bids. The i^{th} bid by bidder n is a K+1-vector $(z_{ni}, v_{ni1},...,v_{niK})$, where the first component (z_{ni}) is a positive quantity and each other component (v_{nik}) is a price or "value." The interpretation of this bid is that bidder n wishes to buy up to z_{ni} units of some version of this product for use in its i^{th} task and will pay up to v_{nik} per unit for version k. In the commodity examples, a task might represent a production process and the version prices may reflect differences in shipping, storage, or processing costs. In the hospital-doctor matching example, each bid may refer to a different role that the hospital needs to fill, each quantity is one, and the value differences may refer to different productivities of a doctor in that role.

If each bidder n makes m_n bids and if the bid prices are interpreted as reported values, then the corresponding equilibrium entails an efficient assignment with respect to the reported values. This assignment is the solution to:

$$\begin{aligned}
&\max_{x_{nik}} \sum_{k=1}^{K}\sum_{n=1}^{N}\sum_{i=1}^{m_n} v_{nik}x_{nik} \text{ subject to} \\
&\sum_{n,i} x_{nik} \le q_k \quad \text{for all product versions } k \ (p_k) \\
&\sum_{k} x_{nik} \le z_{ni} \quad \text{for all bids } n,i \ (\lambda_{ni}) \\
&x_{nik} \ge 0 \qquad \text{for } n,i,k
\end{aligned} \tag{1}$$

Bidder n's goods allocation corresponding to the efficient assignment is $\sum_i x_{ni}$. We use $p_k \ge 0$ to denote the shadow price for the constraint on the total demand for product version k and $\lambda_{ni} \geq 0$ to denote the shadow price for the constraint limiting the quantity assigned to the bid indexed by ni.

The special case of (1) in which for all $n,i,k,$ $q_k = z_{ni} = 1$ is an *assignment problem*: each good is assigned to some use. For such cases, Koopmans and Beckmann (1957) have shown that (1) has an optimal integer solution, that is, an optimal solution with the property that for all $n, i, k,$ $x_{nik} \in \{0,1\}$. It is easy to extend that finding as follows.

<u>Theorem 2</u>. If the parameters q_k and z_{ni} are positive integers, then there is an optimal integer solution of (1), that is, an optimal solution in which each x_{nik} is a non-negative integer.[4]

<u>Proof</u>. We convert (1) to a Koopmans-Beckmann assignment problem by treating each unit of the q_k units of variety k as distinct and assigning it individually. The corresponding problem is:

$$\begin{aligned}
&\max_{\tilde{x}_{nijk}} \sum_{k=1}^{K} \sum_{n=1}^{N} \sum_{i=1}^{m_n} v_{nik} \left(\sum_{j=1}^{q_k} \tilde{x}_{nijk} \right) \text{ subject to} \\
&\sum_{n,i} \tilde{x}_{nijk} \leq 1 \quad \text{for all } k, j \\
&\sum_{k} \tilde{x}_{nijk} \leq 1 \quad \text{for all bids } n,i,j \\
&\tilde{x}_{nijk} \geq 0 \qquad \text{for } n,i,j,k
\end{aligned} \tag{2}$$

[4] A more modern but less elementary argument proceeds by showing that the constraint matrix is *totally unimodular*. For details, see the Wikipedia entry on *unimodular matrix*.

Let $\tilde{x} = (\tilde{x}_{nijk})$ be any integer optimal solution of (2), the existence of which is guaranteed by the Koopmans-Beckmann theorem. Let $x_{nik} \equiv \sum_j \tilde{x}_{nijk}$ Then, *x* is an integer optimal solution to (1). **QED**

Using our definitions of *p* and *λ*, the linear programming dual of (1) is the following:

$$\begin{aligned}
&\min_{p,\lambda} \sum\nolimits_k p_k q_k + \sum\nolimits_{n,i} \lambda_{ni} z_{ni} \text{ subject to} \\
&p_k + \lambda_{ni} \geq v_{nik} \quad \text{for all } n,i,k \\
&p \geq 0, \lambda \geq 0
\end{aligned} \tag{3}$$

In this formulation, for any bid indexed by *ni*, λ_{ni} is zero if the bid is losing bid and is equal to $v_{nik} - p_k$ if the bid is winning, that is, λ_{ni} is the "reported" profit margin.

Given two dual vectors (p,λ) and (p',λ'), we define an ordering in which low goods prices and high bidder profits go together. Thus, define $(p,\lambda) \leq^* (p',\lambda')$ if $p \leq p'$ and $\lambda \geq \lambda'$. In general, the binary operations of *meet* and *join* are specified by $z \wedge z' = \inf\{z,z'\}$ and $z \vee z' = \sup\{z,z'\}$. Using the prescribed ordering, the *meet* $(p'',\lambda'') = (p,\lambda) \wedge (p',\lambda')$ is given by $p''_k = \min(p_k, p'_k)$ and $\lambda''_k = \max(\lambda_k, \lambda'_k)$ and the join $(p''',\lambda''') = (p,\lambda) \wedge (p',\lambda')$ by $p'''_k = \max(p_k, p'_k)$ and $\lambda'''_k = \min(\lambda_k, \lambda'_k)$. A subset of a Euclidean space is a *sublattice* if it is closed under the operation of meet and join.

By inspection, if (p,λ) and (p',λ') both satisfy the constraints in (3), then $(p,\lambda) \wedge (p',\lambda')$ and $(p,\lambda) \vee (p',\lambda')$ satisfy them as well, so the constraint set is a sublattice. An objective *f* is *submodular* if for all *x* and *x′*, $f(x) + f(x') \geq$

$f(x \wedge x') + f(x \vee x')$. This inequality is satisfied by any linear function f because, for linear functions, $f(x) + f(x') = f(x \wedge x') + f(x \vee x')$.

A theorem of Topkis (1978) asserts that the set of minimizers of any submodular function on a sublattice is itself as sublattice and that applies to the optimal solutions of (3). Prices are bounded below by zero and above by maximum valuations for each good. So, the sublattice structure implies that if there are multiple equilibria, then there is one with highest prices and lowest buyer profits and another with the reverse pattern.

We can project these equilibria onto just the space of goods prices to characterize those equilibria. Let π be the optimal value of problems (1) and (3). (These coincide by the duality theorem of linear programming.) By duality, the market-clearing price vectors p for problem (1) are the price vectors that are part of the optimal solutions of (3), as follows:

$$P_0 = \left\{ p \geq 0 \,|\, (\exists \lambda \leq 0) \sum_k p_k q_k + \sum_{n,i} \lambda_{ni} z_{ni} = \pi,\ (\forall n,i,k)\, p_k + \lambda_{ni} \geq v_{nik} \right\}. \tag{4}$$

As we have seen, P_0 is a bounded sublattice. Since the problem (3) is a linear program, P_0 is also closed and convex. Therefore, we have the following.

<u>Theorem 3</u>. The set of market clearing prices P_0 for (1) is the compact, convex sublattice given by (4).

Since P_0 is a compact closed sublattice and bounded below implies that it contains its greatest lower bound, which we denote by $\hat{p}$. The vector $\hat{p}$ and the corresponding profit vector $\hat{\lambda}$ are the unique solution of the following linear program, which selects the lowest-price, highest-profit market-clearing solution.[5]

$$\begin{aligned}
&\min_{p\geq 0,\lambda\leq 0} \sum\nolimits_k p_k - \sum\nolimits_{n,i} \lambda_{ni} \text{ subject to} \\
&\sum\nolimits_k p_k q_k + \sum\nolimits_{n,i} \lambda_{ni} z_{ni} = \pi \\
&p_k + \lambda_{ni} \geq v_{nik} \text{ for all } n,i,k
\end{aligned} \tag{5}$$

Definitions.

1. An *assignment auction* is a mechanism in which bids take the form $(z_{ni}, v_{ni1}, ..., v_{niK})$ consisting of a quantity and a vector of K prices, the allocation is a solution to (1) and the prices lie in the set (4).

2. A *minimum price assignment auction* is an assignment auction in which the prices are $\hat{p}$.

3. The *minimum price assignment auction* μ_F^A is the minimum price assignment auction that selects among multiple optimal solutions in (1) by using the same selection/tie-breaking criterion F as is used for μ_F^{CE}.

4. An *integer assignment auction* is an assignment auction in which the quantity parameters z_{ni} and q_k are restricted to be non-negative integers and the assigned quantities x_{nik} are integers.

Theorem 4. The minimum price assignment auction (μ_F^A) is a tight simplification of the minimal price competitive equilibrium mechanism (μ_F^{CE}).[6]

[5] By the sublattice structure of the constraint set, one can use any objective that is increasing in p and decreasing in λ, that is, any objective that is increasing in the order $\geq^*$.

Proof. In view of Theorem 3, μ_F^A is a simplification of μ_F^{CE}. What remains to prove is that this simplification is tight. Fixing a profile of values v and a profile of simplified strategies s, we must show that if s is not a Nash equilibrium of the complete information game (μ_F^{CE},v), then it is not a Nash equilibrium of the game (μ_F^A,v).

Assume that $s=(s_n,s_{-n})$ is not a Nash equilibrium of (μ_F^{CE},v). Then there is some participant n and strategy $\hat{s}_n$ such that n strictly prefers the resulting outcome: $(x,p)=\omega_F^{CE}(\hat{s}_n,s_{-n})\succ_n \omega_F^{CE}(s)$. By monotonicity and continuity of preferences, for some ε>0 and any p' with $p'_k \le p_k+\varepsilon$ for all k, $(x,p')\succ_n \omega_F^{CE}(s)$.

One profitable deviation $\tilde{s}_n$ for player n in (μ_F^A,v) makes K bids $(z_{ni},v_{ni1},...,v_{niK})$ for $i=1,...,K$, where where $z_{ni}=x_{ni}$, $v_{nii}=p_i+\varepsilon$, and $v_{nik}=0$ for $k\ne i$. With these bids, by construction, if the price vector were set to p, markets would clear and bidder n would buy bundle x_n. Since any market equilibrium is efficient for the simplified reports (and hence total-value maximizing because the simplified reports assert quasi-linear preferences), it follows that n is assigned x_n at every market equilibrium corresponding to the reported strategies. If p' is the selected price vector in the assignment auction, it must support the purchase of x_n, so $p'_k \le p_k+\varepsilon$. So, $\omega_F^A(\tilde{s}_n,s_{-n})=(x,p')=$ $\omega_F^{CE}(\tilde{s}_n,s_{-n})\succ_n \omega_F^{CE}(s)=\omega_F^A(s)$: s is not a Nash equilibrium of (μ_F^A,v). **QED**

[6] The proof actually shows something more. If a simplification of μ_F^{CE} restricts bidders to report quasi-linear preferences and allows them to make the equivalent of the collection of bids $(z_{ni},v_{nii},v_{ni,-i}\equiv 0)$, then the simplification is tight.

An integer assignment auction is a further simplification of an assignment auction, since it restricts the bids that participants are permitted to make. From Theorems 2 and 3, we can also conclude the following.

Theorem 5. A minimum price integer assignment auction exists.

Our emphasis so far has been on simplification, rather than incentives. In view of Theorem 4, results like that Roberts and Postlewaite (1976) about incentives of participants in large markets competitive apply also to assignment auctions.

To understand bidder incentives more thoroughly, observe first that in problem (3), additional bids impose additional lower bound constraints on prices. According to a theorem of Topkis (1978), the result is that the set of optimal solutions "increases" – in particular the component-wise minimum solution becomes higher. So, a bidder who wins items in the assignment auction can never strictly benefit by adding losing bids, since those can only raise prices without affecting the goods assignment. It follows that for a bidder *n* with *unit demand* – *n* has a positive value for at most one item – there is always a best reply that involves submitting one single bid with $z_{n1} = 1$. We argue below that if bidder *n* makes a bid that wins a unit of variety *k*, then it pays the Vickrey price for that unit. Applying the usual Vickrey auction logic, we have the following result.

Theorem 6. In a minimum price integer assignment auction, if participant *n* has unit demand, then it has a dominant strategy, which is to place the single bid $(1, v_{n11}, ..., v_{n1K})$ where each v_{n1k} is its actual values for variety *k*. The price p_k the bidder pays if it wins a unit of variety *k* is the Vickrey price.

Proof. Let $\pi(q)$ denote the maximum value in (1) expressed as a function of the supply vector and let $\pi_{-n}(q)$ denote the maximum value excluding the bids by participant *n*. Since (1) is a linear maximization problem, the minimum shadow prices for the supply constraints are the right-hand derivatives of the value function: $p_k = \partial\pi / \partial^+ q_k$. By Theorem 2, this shadow price can be written as the incremental value of a unit of variety *k*, $p_k = \pi(q_k + 1, q_{-k}) - \pi(q)$.

If bidder *n* acquires x_n, its Vickrey payment is defined to be $\pi_{-n}(q) - \pi_{-n}(q - x_n)$. If *n* bids for a single unit and succeeds in acquiring one – say its goods allocation is x_n consisting of one unit of variety *k*, then $\pi(q) = \pi_{-n}(q - x_n) + v_{n1}(x_n)$ and, since π is convex, $\pi(q + x_n) = \pi_{-n}(q) + v_{n1}(x_n)$. Hence, $\pi_{-n}(q) - \pi_{-n}(q - x_n) = \pi(q + x_n) - \pi(q) =$ $\pi(q_k + 1, q_{-k}) - \pi(q) = p_k$. Whatever bid *n* makes for a single unit, the assignment auction picks the reported value-maximizing allocation and, if it assigns item *k* to *n*, and charges *n* its Vickrey price. So, the standard Vickrey analysis implies that *n*'s best choice among bids for a single unit is always to bid its actual values. We established above that bids for multiple units lead to weakly lower payoffs for *n*. **QED**

IV. Extensions

Two important extensions of the assignment auction allow endogenous supplies and side constraints on demand. The first is motivated partly by the desirability of including multiple sellers and partly by the fact that the simple assignment auction has Nash equilibria with all prices zero. Similar equilibria have been identified in demand-curve and supply-curve auctions when just a single variety is offered for sale, beginning with the path-breaking analysis by Wilson (1979).[7] Low and zero revenue equilibria are eliminated by reserve prices and further benefits can result if the seller adopts upward sloping supply curves. Seller supply offers could be modeled simply by introducing bids with negative quantities, but it is clearer to write the seller's offers separately and to allow multiple sellers in addition to multiple buyers.

An *assignment exchange* supplements bids to buy as described above with offers to sell. Seller *m*'s supply offer is a collection of l_m offer-vectors $(q_{mj}, c_{mj1}, \ldots, c_{mjK})$ where q_{mi} is the quantity offered and c_{mjk} is the minimum price ("cost" or "reserve") for a unit of variety *k*. This is a variation of the *assignment game* of Shapley and Shubik (1972), extended to allow buyers and sellers with multi-unit supplies and demands.

With *M* sellers, the linear program to maximize total value can be written as one of maximizing value minus costs, as follows.

$$
\begin{aligned}
&\max_{x_{nik}, y_{mjk}} \sum_{k=1}^{K} \sum_{n=1}^{N} \sum_{i=1}^{m_n} v_{nik} x_{nik} - \sum_{k=1}^{K} \sum_{m=1}^{M} \sum_{j=1}^{l_m} c_{mjk} y_{mjk} \text{ subject to} \\
&\sum_{n,i} x_{nik} - \sum_{m,j} y_{mjk} \le 0 \quad \text{for all product versions } k \ (p_k) \\
&\sum_{k} x_{nik} \le z_{ni} \qquad \text{for all bids } n, i \ (\lambda_{ni}) \\
&\sum_{k} y_{mjk} \le q_{mj} \qquad \text{for all offers } m, j \ (\theta_{mj}) \\
&x_{nik}, y_{mjk} \ge 0 \qquad \text{for all } m, n, i, j, k
\end{aligned}
\tag{6}
$$

Problem (6) is still a generalized assignment problem, so if the parameter vectors *z* and *q* are integer vectors, then there is an integer optimal solution. The linear programming dual for problem (6) is given by (7) below.

---

[7] See also the important contribution by Klemperer and Meyer (1989) and the discussion in Milgrom (2004). .

$$\begin{aligned}&\min_{p,\lambda,\theta} \sum\nolimits_{m,j} q_{mj}\theta_{mj} + \sum\nolimits_{n,i} z_{ni}\lambda_{ni} \text{ subject to}\\&p_k + \lambda_{ni} \geq v_{nk}\\&p_k - \theta_{mi} \leq c_{mi}\\&p \geq 0, \lambda \geq 0, \theta \geq 0\end{aligned} \tag{7}$$

To analyze this problem, we use an order in which "low" means low prices p, low seller profits θ, and high buyer profits λ: $(p,\theta,\lambda) \leq^* (p',\theta',\lambda')$ if $p \leq p'$, $\theta \leq \theta'$, and $\lambda \geq \lambda'$. Just as for (3), the constraints in (7) define a closed, convex set of dual vectors (p,λ,θ) and the prices and profits are bounded, so the set of equilibrium price vectors P_0 is again a convex, compact sublattice. Its minimum is the market equilibrium with lowest prices, lowest seller profits, and highest bidder profits and its maximum is the equilibrium with highest prices, highest seller profits, and lowest bidder profits.[8]

In the exchange formulation, the simplification is still tight. Indeed, the proof of tightness for the simplified auction given above generalizes to exchanges in a straightforward way.

A second important extension allows buyers or sellers to impose overall limits on the quantities of certain bids that can be accepted. For example, in the medical residency application, an affirmative action program might impose an upper bound on the number of doctors hired with certain characteristics. Or, the FCC may, for policy reasons, wish to limit the number of radio spectrum licenses won by incumbent firms. We incorporate such additional constraints along with the original demand constraints by writing them all together in the form $\sum_{(n,i)\in S}\sum_k x_{nik} \leq z_S$ for some collection of sets S. We require two

[8] Further variations of the exchange formulation are possible and sometimes useful. One restricts the bids for particular products so that only *qualified buyers* can bid for particular products. Another allows participants to act as both buyers and sellers in the same auction/exchange. Combining these two is one way to implement endowed assignment valuations, as discussed further below.

properties. To guarantee that the ordinary demand constraints are represented, we require that the collection $\mathcal{S}$ includes the singleton sets $\{(n,i)\}$. In addition, the collection must form a tree using the set inclusion order: if $S,S' \in \mathcal{S}$ and $S \cap S' \neq \varnothing$, then either $S \subseteq S'$ or $S' \subseteq S$. Writing (1) with these constraints instead of the demand constraints, we get:

$$
\begin{aligned}
&\max_{x_{nik}} \sum_{k=1}^{K}\sum_{n=1}^{N}\sum_{i=1}^{m_n} v_{nik}x_{nik} \text{ subject to} \\
&\sum_{n,i} x_{nik} \le q_k \qquad \text{for all product versions } k \\
&\sum_{(n,i)\in S}\sum_{k} x_{nik} \le z_S \quad \text{for all } S \in \mathcal{S} \\
&x_{nik} \ge 0 \qquad \text{for all } n,i,k
\end{aligned}
\tag{8}
$$

To show that the basic assignment auction can be extended to incorporate side constraints in this way, there are two properties to verify.

Theorem 7. If the right-hand side parameters q_k and z_S are positive integers, then there is an integer optimal solution to problem (8).

Theorem 8. The set of market clearing prices for (8) is a compact, convex sublattice.

Proof of Theorem 7. The proof is by induction on the number κ of additional constraints, that is, sets in $\mathcal{S}$ that are not singleton sets $\{n,i\}$. When there are no additional constraints ($\kappa = 0$), problem (8) is the standard assignment auction problem, so an integer solution exists. Suppose inductively that there is an integer optimal solution to all assignment problems with at most κ additional constraints and suppose the actual problem has $\kappa+1$ such constraints. Let $\hat{S}$ be a constraint set whose only subsets in the collection $\mathcal{S}$ are the singleton subsets. The collection $\hat{\mathcal{S}} = \mathcal{S} - \{\hat{S}\}$ has κ elements and a hierarchical structure. Let M be any strict upper bound for the set of values (for example, $M = 1 + \max_{nik} v_{nik}$).

We construct a new problem with $K+1$ products, with values v_{nik} for varieties $k = 1,...,K$ the same as in (8) and $v_{ni,K+1} = M$ for $(n,i) \in \hat{S}$ and $v_{ni,K+1} = 0$ otherwise. Let the supply of good K+1 be $q_{K+1} = -z_{\hat{S}} + \sum_{(n,i)\in\hat{S}} z_{ni}$. The new problem is an assignment problem with κ additional constraints, as follows.

$$\begin{aligned}
&\max_{x_{nik}} \sum_{k=1}^{K+1} \sum_{n=1}^{N} \sum_{i=1}^{m_n} v_{nik} x_{nik} \text{ subject to} \\
&\sum_{n,i} x_{nik} \le q_k \qquad \text{for all product versions } k = 1,...,K+1 \\
&\sum_{(n,i)\in S} \sum_{k=1}^{K+1} x_{nik} \le z_S \quad \text{for all } S \in \hat{\mathcal{S}} \\
&x_{nik} \ge 0 \qquad \text{for all } n,i,k
\end{aligned} \tag{9}$$

$$\begin{aligned}
&= Mq_{K+1} + \max_{x_{nik}} \sum_{k=1}^{K} \sum_{n=1}^{N} \sum_{i=1}^{m_n} v_{nik} x_{nik} \text{ subject to} \\
&\sum_{n,i} x_{nik} \le q_k \qquad \text{for all product versions } k = 1,...,K \\
&\sum_{(n,i)\in S} \sum_{k=1}^{K} x_{nik} \le z_S \quad \text{for all } S \in \mathcal{S} \\
&x_{nik} \ge 0 \qquad \text{for all } n,i,k
\end{aligned} \tag{10}$$

Due to the high coefficient M, $\sum_{(n,i)\in\hat{S}} x_{ni,K+1} = \sum_{n,i} x_{ni,K+1} = q_{K+1}$, where $q_{K+1} = -z_{\hat{S}} + \sum_{(n,i)\in\hat{S}} z_{ni}$. Summing the constraints $\sum_{k=1}^{K+1} x_{nik} \le z_{ni}$ over $(n,i) \in \hat{S}$ and subtracting q_{K+1} from both sides leads to $\sum_{(n,i)\in\hat{S}} \sum_{k=1}^{K} x_{nik} \le z_{\hat{S}}$, yielding the equality between problems (9) and (10). By the inductive hypothesis, the first problem (which is an assignment problem with κ additional constraints) has an integer optimal solution, and that is also an optimal solution of (10). Since (10) incorporates the problem (8) as a sub-problem, the relevant part of the integer optimal solution is also an optimal solution for (8). **QED**

Proof of Theorem 8. For any set S except the root $\bar{S} = \bigcup_{S \in \mathcal{S}} S$, denote the immediate predecessor by $S^{+} = \bigcap_{S' \supset S} S'$. The linear programming dual for (8) is the following.

$$\begin{aligned} &\min_{p,\lambda} \sum\nolimits_{k} p_k q_k + \sum\nolimits_{S \in \mathcal{S}} \lambda_S z_S \text{ subject to} \\ &p_k + \sum\nolimits_{\{S \mid (n,i) \in S\}} \lambda_S \geq v_{nik} \quad \text{for all } n,i,k \\ &p \geq 0, \lambda \geq 0 \end{aligned} \tag{11}$$

Define $\hat{\lambda}_S = \sum\nolimits_{S' \supseteq S} \lambda_{S'}$ and $\lambda_S = \hat{\lambda}_{S^+} - \hat{\lambda}_S$ and rewrite (11) as follows.

$$\begin{aligned} &\min_{p,\hat{\lambda}} \sum\nolimits_{k} p_k q_k + \sum\nolimits_{S \in \mathcal{S} - \{\bar{S}\}} (\hat{\lambda}_S - \hat{\lambda}_{S^+}) z_S + \hat{\lambda}_{\bar{S}} z_{\bar{S}} \text{ subject to} \\ &p_k + \hat{\lambda}_{ni} \geq v_{nik} \quad \text{for all } n,i,k \\ &\hat{\lambda}_S - \hat{\lambda}_{S^+} \geq 0 \quad \text{for all } S \in \mathcal{S} - \{\bar{S}\} \\ &p \geq 0, \hat{\lambda}_{\bar{S}} \geq 0 \end{aligned} \tag{12}$$

Define an order by $(p,\hat{\lambda}) \leq^{*} (p',\hat{\lambda}')$ if $p \leq p'$ and $\hat{\lambda} \geq \hat{\lambda}'$. By inspection, the constraint set is a sublattice and, by Topkis's Theorem, the set of optimal solutions is a sublattice as well. It follows that the projection onto the space of goods prices is similarly a convex, compact sublattice. **QED**

V. Discussion

The assignment auction has been constructed to be far simpler than the full direct mechanism μ_F^{CE}, employing a form of bid that is usable and similar to existing forms, yet which is still flexible enough to express many one-to-one substitution possibilities.

One important topic to investigate is the scope of the set of bids. Is it large enough to express the most common one-for-one substitute preferences? For example, suppose that a certain buyer can process ten units of grain today and another ten units tomorrow and has capacity to store four units from today until tomorrow at a cost of *c* per unit stored. If the value of units of v_1 today and v_2 tomorrow, then it can express these values in an assignment auction by making three bids: $(10, v_1, 0)$, $(6, 0, v_2)$ and $(4, v_2 - c, v_2)$. These bids correspond intuitively to three categories: units today, units tomorrow, and units for tomorrow that could also be acquired either today or tomorrow.[9]

One interesting possibility that is missing from this demand specification is that participant *n* may already have some units in its inventory that it could use today or tomorrow by incurring some opportunity cost. In that case, its values for product are not described by the simple *assignment valuations* used above but it can be described by an *endowed assignment valuation* (Hatfield and Milgrom (2005)). This kind of demand could, in principle, be accommodated in the preceding example by using the exchange design in the following way. Introduce a third variety of the product ("units in bidder *n*'s inventory"). Suppose, for example, that *n* has five units already in inventory with an opportunity value of c_3 and that those units could be used today or tomorrow. Then, its

[9] There is a close similarity between the multiple bids used here and the multiple lists used by some hospitals in the NRMP. If a hospital needs at least one doctor with particular skills, or at least some number of women, it can partially encode that by creating a separate list for the highly demanded group. .

demand could be expressed by n offering to sell $(5,0,0,c_3)$ to qualified buyers and altering the other bids to $(10,v_1,0,v_1)$, $(6,0,v_2,v_2)$ and $(4,v_2-c,v_2,v_2)$. Only bidder n would be qualified to bid for this "private" variety. The resulting demand accurately incorporates n's existing inventory values into its expressed demands for the auctioned goods. Because of the extra complexity of expression that is entailed by endowed assignment valuations, however, attention here has been focused on simple assignment valuations.

A frequently helpful way to assess the complexity of a new mechanism design such as the simple assignment auction is by studying its relationship to existing designs of established practicality. One design closely related to the assignment auction is the sealed-bid uniform price auction for a single commodity. In that auction, the auctioneer constructs a demand curve from a set of price-quantity bids and crosses that curve that with a fixed supply to determine a clearing price. If there were but a single variety of the good, the assignment auction would reduce to precisely this sort of a sealed-bid uniform price auction. If there are multiple varieties of the good, bidders could continue to bid in the familiar way specifying a price-quantity pair for each good, but the assignment auction allows them the additional option of expressing a substitution possibility, depending on the prices of the several varieties.

Something resembling the minimum price assignment auction design could also emerge in steps from some ascending clock auctions. For example, a bid of $(10,v_1,v_2)$ in the assignment auction plays much the same role as an instruction to an agent in an ascending clock auction which says "buy 10 units for prices up to my values (v_1,v_2) and always bid for whichever is the better value for me." If all bidders were to give instruction of this form to the auctioneer, then the result would be almost the same as in an assignment auction.

The comparison between these designs, however, is inexact. Ascending clock auctions enable bidders to follow a wider range of strategies than in the assignment auction, which is, after all, a simplification! This potential disadvantage of simplification is offset by two other factors. First, bidding in an ascending clock auction need not be consistent with any demand function at all. Second, clock auctions operate in real time, which forces a tradeoff between auction duration and accuracy. If bid increments in a clock auction are large, then the auction closes quickly but market clearing is subject to approximation error. When increments are small, multi-product ascending auctions can take many hours to run (the spectrum auctions sometimes take months to reach completion).[10] Auctions that run over several hours can be especially costly in international commodity trading, with potential buyers sometimes located on opposite sides of the earth. Time is money: a sealed bid design can save a great deal of human time and effort.

The most *ad hoc* part of this analysis has been the judgment about useful principles for simplification. We have employed the idea of "tight simplifications" here to guide our judgment, but any concept for mechanism design applications that relies on full information pure Nash equilibrium is inherently suspect. This problem is a difficult

[10] Real-time clock auctions suffer from approximation error because the auctioneer must guess at bid increments. To be completed quickly enough to be practical, bid increments are often set large enough to risk overshooting market clearing prices. Gul and Stacchetti (2000) abstract from that the related problems of auction duration and increment size by assuming that all values are integers and that bid increments are of size one.

Approximation error can be mitigated by clever design, for example by the introduction of "intra-round bidding" as suggested by Ausubel (2007), but there is some irreducible approximation error associated with any clock auction, because all require that the auctioneer must name parts of the price trajectory before bidders report their demands.

one. The sheer complexity of the environments studied here, in which bidders may find it too costly even to formulate a complete set of preferences, makes it equally unconvincing to rely on concepts embedding high levels of Bayesian rationality.

In the world of practical design, conceptual difficulties like this one are no cause for inaction. We proceed with the principles we have, while hoping that better grounded principles can eventually be used to guide simplification in mechanism design.

References

Ausubel, Lawrence. 2004. "An Efficient Ascending Bid Auction for Multiple Objects." *American Economic Review*, 94:5, pp. 1452-75.

Ausubel, Lawrence. 2007. "System and Method for an Efficient Dynamic Multi-Unit Auction." *US Patent Number 7,165,046*.

Crawford, Vincent P. 2004. "The Flexible-Salary Match: A Proposal to Increase the Salary Flexibility of the National Resident Matching Program." *UCSD Working Paper*.

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.

Gul, Faruk and Ennio Stacchetti. 2000. "The English Auction with Differentiated Commodities." *Journal of Economic Theory*, 92:1, pp. 66-95.

Hatfield, John and Paul Milgrom. 2005. "Matching with Contracts." *American Economic Review*, 95:4, pp. 913-35.

Kelso, Alexander and Vincent Crawford. 1982. "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica*, 50, pp. 1483-504.

Klemperer, Paul and Margaret Meyer. 1989. "Supply Function Equilibria in Oligopoly Under Uncertainty." *Econometrica*, 57:6, pp. 1243-77.

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

Milgrom, Paul. 2000. "Putting Auctions Theory to Work: The Simultaneous Ascending Auction." *Journal of Political Economy*, 108:2, pp. 245-72.

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

Milgrom, Paul. 2007. "Simplified Mechanisms with Applications to Sponsored Search and Package Bidding." *Working paper, Stanford University*.

Roberts, John and Andrew Postlewaite. 1976. "The Incentives for Price-Taking Behavior in Large Exchange Economies." *Econometrica*, 44:1, pp. 115-29.

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, pp. 748-80.

Shapley, L. and M. Shubik. 1972. "The assignment game I: The core." *International Journal of Game Theory*, 1, pp. 111-30.

Topkis, Donald. 1978. "Minimizing a Submodular Function on a Lattice." *Operations Research*, 26:2, pp. 305-21.

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

Wilson, Robert. 1979. "Auctions of Shares." *Quarterly Journal of Economics*, XCIII:4, pp. 675-89.