Research Problems in Combinatorial Auctions

Rakesh V. Vohra*

October 29, 2001

Abstract

Many auctions involve the sale of a variety of distinct assets. Examples are airport time slots, delivery routes, network routing and furniture. Because of complementarities or substitution effects between the different assets, bidders have preferences not just for particular items but for sets of items. For this reason, economic efficiency is enhanced if bidders are allowed to bid on bundles or combinations of different assets. This note outlines some of the underlying theory and identifies some research questions. The reader interested in more details on research in this are should consult the survey paper by de Vries and myself on combinatorial auctions. I am greatful for useful conversations with Sushil Bikhchandani, Sven de Vries, Peter Eso, Tuomas Sandholm, James Schummer and Subhash Suri.

(*Auctions; Combinatorial Optimization*)

*Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston IL 60208. e-mail: r-vohra@nwu.edu

1 Introduction

Many auctions involve the sale of a variety of distinct assets. Examples are the FCC spectrum auction (`http://www.fcc.gov/wtb/auctions/`) and auctions for airport time slots (Rassenti *et al.* 1982) and delivery routes (Olson *et al.* (200), Caplice 1996). Because of complementarities or substitution effects between different assets, bidders have preferences not just for particular items but for sets of items, sometimes called bundles.

To illustrate, suppose you must auction off a dining room set consisting of four chairs and a table. Would you wish to auction off the entire set or run five separate auctions for each piece? The answer depends, of course, on what bidders care about. If every bidder is interested in the dining room set and nothing less, the first option is preferable. If some bidders are interested in the set but others are interested only in a chair or two it is not obvious what to do. If you believe that you can raise more by selling off the chairs separately than the set, the second option is preferable. Notice, deciding requires a knowledge of just how much bidders value different parts of the ensemble. For this reason, economic efficiency is enhanced if bidders are allowed to bid directly on *combinations* of different assets instead of bidding only on individual items. Auctions where bidders are allowed to submit bids on combinations of items are usually called combinatorial auctions. 'Combinational auctions' is more accurate, but in this survey we will comply with convention.

2 The Set Up

Let N be the set of bidders ($n = |N|$) and M the set of distinct objects. Each bidder, $j \in N$ is endowed with a **value function**, v^j which specifies the monetary value (non-negative) that bidder j assigns to every subset of objects. Thus, the value that bidder j assigns to consuming $S \subseteq M$ is $v^j(S)$. These value functions are private information to each bidder. Notice that bidder j cares only about what she consumes and not what others consume. So, we are in the private values, no externality world. This restriction holds throughout. The realm of interdependent values has even more problems and I direct the reader to Jehiel and Moldovanu (2001) for some of the issues.

Each bidder assumes that the other bidders value functions are independent draws from a finite set V of value functions with commonly known distribution F. The probability that agent j has value function u will be denoted $F(u)$. The probability that the n-tuple $\mathbf{v} = (v^1, \ldots, v^n) \in V^n$ is realized is denoted $\Pi_{j=1}^n F(v^j)$. For convenience this will be abbreviated to $F(\mathbf{v})$.

We assume an auctioneer interested in auctioning off the elements of M so as to maximize expected revenue. By the revelation principle [see Myerson (1981)] we can restrict attention to a direct revelation scheme. In such a scheme the auctioneer announces how he will allocate the objects amongst the bidders and the payments he will extract from each as a function of the announced value functions. Then bidders are asked to announce their value functions. The auctioneer is constrained to choose an allocation procedure and payment function that satisfy two constraints:

1. Induce bidders to reveal their actual valuations (**incentive compatibility**).

2. No bidder is made worse off (in expectation) by participating in the auction.

To define an **allocation**, let $y(S, j) = 1$ if the bundle $S \subseteq M$ is allocated to $j \in N$ and zero otherwise. An allocation is a solution to:

$$\sum_{S \ni i} \sum_{j \in N} y(S, j) \leq 1 \ \forall i \in M$$

$$\sum_{S \subseteq M} y(S, j) \leq 1 \ \forall j \in N$$

$$y(S, j) = 0, 1 \ \forall S \subseteq M, \ j \in N$$

The first constraint ensures that overlapping sets of goods are never assigned. The second ensures that no bidder receives more than one subset.

An **allocation rule** is a mapping from an n-tuple of value functions to an integer solution to the system above. If A is an allocation rule and $\mathbf{v} = (v^1, v^2, \ldots, v^n)$ an n-tuple of value functions, we will write $A(\mathbf{v})$ to mean the allocation selected by the rule A.

A **payment rule** is a mapping from an n-tuple of value functions to an n-tuple of payments, one for each bidder. If P is a payment rule and $\mathbf{v} = (v^1, v^2, \ldots, v^n)$ an n-tuple of value functions, we will write $P_j(\mathbf{v})$ to mean the payment that bidder j must make.

3 Revenue Maximizing Auction

Here I will formulate as a mathematical program the problem of finding the revenue maximzing auction. This is just a a specialization of the usual formulation for optimal auctions with multidimensional types. The survey paper by Rochet and Stole (2000) has more details on this subject (see also Jehiel *et al.* (1999)). No closed form solution to the problem of optimal auctions with multidimensional types is known and it is unlikely that any exists. Under fairly severe restrictions on value functions and the space of types (amounting to the types being one dimensional) results are available, see Rochet and Stole (2000).

There is one decision variable that represents the choice of allocation rule and another to represent the choice of payment rule. The objective is

$$\max_{A,P} \sum_{\mathbf{v} \in V^n} F(\mathbf{v})[\sum_{j \in N} P_j(\mathbf{v})].$$

Incentive compatibility requires that for each $j \in N$ with value function v^j and all $u \neq v^j$:

$$\sum_{\mathbf{v^{-j}} \in V^{n-1}} [v^j(A(v^j, \mathbf{v^{-j}})) - P_j(v^j, \mathbf{v^{-j}})]F(\mathbf{v^{-j}}) \geq \sum_{\mathbf{v^{-j}} \in V^{n-1}} [v^j(A(u, \mathbf{v^{-j}})) - P_j(u, \mathbf{v^{-j}})]F(\mathbf{v^{-j}}).$$

That is, the expected utility from truthfully reporting ones value function should be at least as large as the expected utility from reporting some other value function. Given incentive compatibility we can write the individual rationality constraint as

$$\sum_{\mathbf{v^{-j}} \in V^{n-1}} [v^j(A(v^j, \mathbf{v^{-j}})) - P_j(v^j, \mathbf{v^{-j}})]F(\mathbf{v^{-j}}) \geq 0$$

for all $j \in N$ and value functions v^j.

How might one solve this? First fix a choice for allocation rule, A. Let

$$\nu^j(A, u) = \sum_{\mathbf{v^{-j}} \in V^{n-1}} v^j(A(u, \mathbf{v^{-j}}))F(\mathbf{v^{-j}})$$

and

$$\rho^j(u) = \sum_{\mathbf{v^{-j}} \in V^{n-1}} P_j(u, \mathbf{v^{-j}})F(\mathbf{v^{-j}}).$$

Then we can rewrite the incentive compatibility and individual rationality constraints as follows:

$$\rho^j(v^j) - \rho^j(u) \le \nu^j(A, v^j) - \nu^j(A, u),$$

and

$$\rho^j(v^j) \le \nu^j(A, v^j).$$

The objective function can be written as:

$$\sum_{\mathbf{v}\in V^n} F(\mathbf{v})[\sum_{j\in N} P_j(\mathbf{v})] = \sum_{j\in N}\sum_{v\in V} F(v)\rho^j(u).$$

Thus the optimization problem becomes:

$$\max \sum_{j\in N}\sum_{v\in V} F(v)\rho^j(v)$$

$$\text{s.t. } \rho^j(v) - \rho^j(u) \le \nu^j(A, v) - \nu^j(A, u),\ \forall v, u \in V$$

$$\rho^j(v) \le \nu^j(A, v),\ \forall v \in V.$$

Since the allocation rule A is fixed, the only variables are the ρ's. There are at most two of them per constraint with coefficients of +1 and -1. Hence, given A, the problem of finding a payment rule that enforces incentive compatibility and individual rationality is a network flow problem (or more precisely the dual to one).[1] For each possible choice of allocation rule A, we can determine via the duality theorem of linear programming whether a payment rule that is incentive compatible exists. Here is a recipe for determining the revenue maximizing auction. Choose an allocation rule. Solve the shortest path problem induced by this allocation rule to determine the revenue maximizing payments that are incentive compatible and individually rational. Choose another allocation rule and repeat.

- For what classes of valuation functions can we solve the optimal auction problem, either in closed form or with a polytime algorithm?
- Can one find the revenue maximizing auction from amongst a restricted class of allocation rules?
- For one dimensional types, a Vickrey auction with a reserve price is revenue maximizing. For the general case, how close to optimal is a Vickrey auction with reserve prices?

[1] Introduce one vertex for each value function in V. For each ordered pair (v, u), introduce an arc directed from v to u with length $\nu^j(A, v), -\nu^j(A, u)$. The dual problem is to find the shortest path tree rooted at the source.

4 Efficient Auctions

One very popular restriction of allocation rules, is the rule that picks the **efficient allcoation**. The efficient allocation is the solution to the following problem (CAP):

$$\max \sum_{j \in N} \sum_{S \subseteq M} v^j(S) y(S, j)$$

$$\text{s.t.} \sum_{S \ni i} \sum_{j \in N} y(S, j) \leq 1 \ \forall i \in M$$

$$\sum_{S \subseteq M} y(S, j) \leq 1 \ \forall j \in N$$

$$y(S, j) = 0, 1 \ \forall S \subseteq M, \ j \in N$$

It is an instance of a set packing problem. In the context of combinatorial auctions this is often called the winner determination problem. Winner determination is a well worked area. See Rothkopf *et al.* (1998), Sandholm *et al.* (2001) and de Vries and Vohra (2001) for a discussion of this subject.

4.1 A Comprehensive Formulation

I have given one formulation for winner determination, but it is not general enough to encompass the formulations that have been used in special cases.

$$\max \sum_{i \in N} \sum_{S \subset \Omega_i} v^j(S) y(S, j)$$

$$\text{s.t.} \sum_{S \ni j} y(S, j) \leq 1 \ \forall j \in N \, \forall S \subset \Omega_i \quad GCAP_1$$

y$^j \in P_j^A \ \forall j \in M \quad GCAP_2$
y $\in P^A \quad GCAP_3$
y$^j \in P_j^B \ \forall j \in M \quad GCAP_4$
y(S,j) = 0,1 $\quad \forall j \in M \, \forall S \subset M \quad GCAP_5$

Here, P_j^A and P_j^B and P^A are polyhedral sets.

The sets $\Omega_i \subseteq 2^M$ model restrictions on what bidders on can bid. They can be fixed by the auctioneer or she might permit bidders to construct them themselves (subject to some constraints as in FCC-restrictions for auction #31 on the number of packages). The constraints c-gcap-1 ensure that no item is allocated to more than one bidder. The constraints c-gcap-2 are imposed by the auctioneer and enforce capacity constraints on the bidders; for example no bidder is supposed to win more than two items, no bidder is supposed to win more than 40% of the total business, etc. In fact, the auctioneer could construct polytopes P_j^A that in effect restrict Ω_j. To avoid this we require $P_j^A \cap \mathbf{N}^{\Omega_j}$ to be full dimensional.

Constraints c-gcap-3 permit the auctioneer to restrict the overall allocation. For example, the allocation must be edges that form a path or a tree. We assume that P^A can not be described as the cartesian product of an interval on the coordinate axis and a lower-dimensional polytope.

Constraints c-gcap-4 allow each bidder to restrict the allocations he might win. For example, if he has a subadditive valuation, he might put $P_i^B = \{y \in \mathbf{R}^{\Omega_j} | \sum_{S\in\Omega_j} y(S, j \leq 1\}$ to ensure that he doesn't pay more than what he bid.

Finally, c-gcap-5 ensures that we end with an integral allocation.

4.2 Vickrey Auction

The incentive compatible and individually rational payment rule consistent with the efficient allocation that maximizes the auctioneer's revenue is the (expected) Vickrey payment rule.[2] The argument is simple and the reader is referred to Williams (1999) or Krishna and Perry (1998) for the details. One consequence of this derivation is that the efficient auction that maximizes revenue for the seller (subject to individual rationality) is the Vickrey auction.

To describe the sealed bid Vickrey auction, define for any $K \subseteq N$ let

$$V(K) = \max \sum_{j\in K} \sum_{S\subseteq M} v^j(S) y(S, j)$$

$$\text{s.t. } \sum_{S\ni i} \sum_{j\in N} y(S, j) \leq 1 \ \forall i \in M$$

$$\sum_{S\subseteq M} y(S, j) \leq 1 \ \forall j \in N$$

[2]One does not get Vickrey directly but an an expected value version because we require only (Bayesian) incentive compatibility rather than dominance. However a similar argument can be used to produce the Vickrey auction. Also, to save on space what I am calling the Vickrey auction should be called the generalized vickrey auction or the Vickrey-Clarke-Groves auction.

$$y(S,j) = 0,1 \; \forall S \subseteq M, \; j \in N$$

The auction works as follows:

1. Agent j reports v^j. There is nothing to prevent agent j from misrepresenting themselves. However, given the rules of the auction, it is a weakly dominant strategy to bid truthfully.

2. The seller implements the efficient allocation, i.e., the one that solves:

$$V = \max \sum_{j \in N} \sum_{S \subseteq M} v^j(S) y(S,j)$$

$$\text{s.t.} \sum_{S \ni i} \sum_{j \in N} y(S,j) \leq 1 \; \forall i \in M$$

$$\sum_{S \subseteq M} y(S,j) \leq 1 \; \forall j \in N$$

$$y(S,j) = 0,1 \; \forall S \subseteq M, \; j \in N$$

   Call the efficient allocation y^*

3. The seller also computes $V(N \setminus j)$ for each $j \in N$.

4. Each bidder $j \in N$ is charged

$$V(N \setminus j) - [V - \sum_{S \subseteq M} v^j(S) y^*(S,j)].$$

   Thus bidder k's payment is the difference in 'welfare' of the other bidders without him and the welfare of others when he is included in the allocation. Notice that the payment made by each bidder to the auctioneer is non-negative.

Under this auction each bidder j obtains a profit of $V(N) - V(N \setminus j)$. This last term is sometimes called bidder j's **marginal product**. Furthermore it is a dominant strategy for each bidder to bid truthfully. Executing the Vickrey auction requires the solution to $n+1$ optimization problems.

- Suppose agents preferences are such that the agents are subtitutes condition holds (i.e. $V(\cdot)$ is submodular). Is the Winner determination problem in this case polynomial in the number of objects and bidders?

This is the case for all known cases where the agents are substitutes property holds.[3]

- In some special cases [Hershberger and Suri (2001) and Bikhchandani *et al.* (2001), de Vries *et al.* (2001)] this can be reduced to at most a single optimization problem of the same complexity as that of finding the efficient allocation. What other cases are there?

- Is there any reason to suppose that auction problems give rise to set packing problems with a very different structure from those already encountered elsewhere? If the answer is no, one can rely on existing set packing solvers. Leyton-Brown *et al.* (2000) are collecting a test suite of winner determination problems and are on the hunt for 'real' problem instances.

- Winner determination with side constraints is also of interest, see Davenport and Kalaganam (2001) as well as Sandholm and Suri (2001).

- If the problem of finding an efficient allocation is NP-hard, we may have to rely on an approximation algorithm. An approximation algorithm is an allocation rule. Which approximation algorithms can be supported by an incentive compatible payment rule? Lehmann *et al.* (2000), for example, uses the greedy allocation and the Vickrey payment rule. In general, this is not incentive compatible, but in a very special case of single minded bidders it is.[4]

- What is the trade-off between revenue and efficiency when implementing an approximately efficient auction that can be supported by an incentive compatible payment rule?

- Executing the Vickrey auction requires that each bidder identify and transmit a value for each combination of items. How is this to be done in a computationally feasible way? Nisan (1999) considers the use of a bidding language and given a language for asks what preferences over subsets of objects can be correctly represented by the language. Nisan (2001) asks what is the number of subsets of M that the auctioneer

[3]For example, diminshing marginal utilities, heterogenous objects with unit demand and gross substitutes.

[4]For one dimensional types there is a literature on which allocation rules are supported by incentive compatible payment rule, see pg. 257 of Fudenberg and Tirole (1992).

needs values for in order to determine the efficient allocation? Nisan shows that values for $O(2^{|M|})$ subsets (in the worst case) are needed. For what preference structures can one do better? Or, perhaps one should be asking different kinds of quetsions to elicit valuation information (Conen and Sandholm 2001).

4.3 Complements and Substitutes

The Vickrey auction gets a lost of stick for being counter-intuitive. One can construct examples where as the number of bidders increase the revenue to the seller drops. Also, in procurement settings, one can construct examples where the Vickrey auction is overly generous, i.e., paying astounding multiples of what something is worth. I want to take some time to stamp out the view that these are counter-intuitive phenomena.

The key to understanding the behavior of the Vickrey auction is the agents are substitutes condition. In the selling context, recall, this means that $V(N)$, the value of the efficient allocation is submodular. In the procurement context where the cost of the efficient allocation ($C(N)$) is relevant, we require that $C(N)$ be supermodular. In both cases an economic interpretation of the condition is not obvious, so I prefer an alternative formulation.

Let $MP(S)$ denote the marginal product of the subset S of agents. Thus $MP(S) = V(N) - V(N \setminus S)$ or $MP(S) = C(N \setminus S) - C(N)$ depending on whether the auctioneer is selling or buying. The substitutes condition is then

$$MP(S) \geq \sum_{j \in S} MP(j) \ \forall S \subseteq N.$$

The marginal product of a group is at least as large as the sum of the individual marginal products.

Now I would like you to consider two problems defined on graphs. First fix a connected graph G with edge set E and vertex set V. Every edge in the graph is owned by a single agent and no agent owns more than one edge.[5] The length/cost of an edge is private information to the agent who owns the edge. The first problem is the minimum cost spanning tree problem. If A is any subset of agents, denote by $M(A)$ the length of the minimum spanning tree using edges in A alone.[6] In Bikhchandani *et al.* (2001) it is shown that $M(A)$ is supermodular.

[5] For this reason I will use the words edge and agent interchangeably.

[6] If A does not span the graph, take $M(A) = \infty$.

The second problem is the shortest path problem. Designate a source, s and sink t and let $L(A)$ be the length of the shortest $s-t$ path using edges only in A.[7] Also to make things interesting, assume that the minimum $s-t$ cut is at least two. As shown in Bikhchandani *et al.* (2001) $L(A)$ is not in general supermodular.

In the first case the auctioneer would like to buy the edges necessary to build the minimum spanning tree. If the auctioneer uses a Vickrey auction, what does she end up paying for an edge, e, say, in the MST? If T is the MST then she buys edge e for the cost of the cheapest edge f that forms a cycle with e in $T \cup f$. Thus, if the number of edges in G were to increases, the price that the auctioneer would have to pay for an edge can only go down. Second, if the ratio of the most expensive edge to the cheapest were α, say, each agent in T would receive at most α times their cost. In short everything about the Vickrey auction in this case conforms to our inuition about how auctions should behave. Why? This is because agents are substitutes. An agent is included in the MST if it has low cost and there is no other agent with with a lower cost that will perform the same 'function'. That function is to 'cover a cut'. Thus no agent depends on another to make it into the MST, but another agent could crowd it out of the MST. This pits the agents against each other an benefits the auctioneer. Alternatively, the agents benefit by banding together, this in effect is what the substitutes property means.

Now turn to the shortest path problem. Here the agents are substitutes property does not hold. In fact, the reverse is almost true agents complement each other. Whether an agent makes it into the shortest $s-t$ path depends on whether other agents make it in. To see why suppose G consists of just two disjoint $s-t$ paths each involving K edges. The edges in the 'top' path have cost 1 each while in the 'bottom' path they have cost $1+r$ each. The Vickery auction applied to this case would choose the top path. Each agent on the top path has a marginal product of rK and so receives a payment of $1+rK$. The total payment by the auctioneer is $K+rK^2$, which, for appropriate r and K is many times larger than the length of the shortest path. A clear case of the 'generosity' of the Vickrey auction. Furthermore, if we increase the number of agents by subdividing a top edge and assigning each segment a cost of $1/2$ the payments increase without an increase in the length of the path! This 'odd' behavior occurs because the agents complement each other. No agent on the top path can be selected unless all other agents on the top

[7] If A does not contain a $s-t$ path set $L(A) = \infty$.

path are selected. Agents on the same path are not in competition with each other. The marginal product of the agents on the top path is less than the sum of their individual marginal products.

When there are such severe complementarities and our goal is to reduce the payments made by the auctioneer, will some other incentive compatible mechanism work? This is a line taken by Archer and Tardos (2002) who suggests that one relax the efficiency requirement. In the shortest path context this would mean accepting a path that is perhaps only slightly longer than the shortest path but the auctioneer ends up paying far less for it. Archer and Tardos (2002) replaces efficiency by a collection of 'monotonicty' type conditions and shows that the same 'odd' behavior persists.[8] In light of the complementary nature of agents this is not surprising. In fact I think it is the complementary nature of agents coupled with incentive compatibility that drives the generosity of the Vickrey auction rather than the desire for efficiency.

To see why, return to the simple graph with two disjoint $s-t$ paths. Now that we understand the complementary nature of the agents, we recognize that the auctioneer wants $s-t$ paths to compete with each other rather than edges. How can this be arranged? One thing the auctioneer could is simply tell the agents on the top path to band together and submit a single bid and the ones on the bottom a single bid. This forces competition between the two paths and should lower payments. Now think about the winning group of bidders. They receive a a certain amount of money to divide between them. How do they do it? The winning group of bidders face a classical public goods problem. The only way to induce each agent to truthfully reveal his cost is to throw money at them. If the auctioneer does not provide the money, there is no incentive for the bidders to band together.

4.3.1 False Name Bidding and Bribe-Proofness

When agents are substitiutes, the Vickrey auction is vulnerable to manipulation by bidders banding together. When the substitutes condition is not satisfied, then agents can benefit by 'splitting' up. A single agent masquerades as many different agents. This is called fasle name bidding. The shortest path example illustrates this very nicely. Take any agent on the top path and have her split into two edges of length $1/2$ each.

[8] One can think of these additional conditions as being properties of a path that would be generated by optimizing some 'nice' function on paths other than just length.

Makoto Yoko has proposed some mechanisms that would be invulnerbale to false name bidding. It turn out that false name bidding is related to 'bribe-proofness' a concept studied in Schummer (2001). A mechanism is bribe proof if no agent can pay another to lie such that both a better off than. One can think of false name bidding as introducing a new agent who can be bribed by an existing agent. Under a standard domain richness condition, Schummer (2001) that a bribe proof rule must be a constant rule.

- When agents are substitutes one should be able to 'bound' the generosity of the Vickrey auction in a nice way. In fact it should be possible to give up on efficiency and reduce payments.
- When agents are not substitutes, just how complementary are they? How should one package the bidders together so as to stimulate competition and at the same time reduce the public goods problem that each package of bidders face?

5 Indirect Mechanism's

A direct mechanism places a huge computational burden upon the auctioneer and imposes a considerable communication cost between bidder and auctioneer. For this (and other reasons) there has been interest in indirect mechanisms. In the context of auctions such mechanisms are referred to as ascending/iterative auctions. Purists who know the latin root of the word auction, carp that the qualifier 'ascending' is redundant. For this and other reasons, I will use the term iterative auction.

Iterative auctions come in two varieties (with hybrids possible). In the first, bidders submit, in each round, prices on various allocations. The auctioneer makes a provisional allocation of the items that depends on the submitted prices. Bidders are allowed to adjust their price offers from the previous rounds and the auction continues. Such auctions come equipped with rules to ensure rapid progress and encourage competition. Iterative auctions of this type seem to be most prevalent in practice.

In the second type, the auctioneer sets the price and bidders announce which bundles they want at the posted prices. The auctioneer observes the requests and adjusts the prices. The price adjustment is usually governed by the need to balance demand with supply.

Call auctions of the first type *quantity setting*, because the auctioneer sets the allocation or quantity in response to the prices/bids set by bidders. Call the the second *price setting* because the auctioneer sets the price. Quantity setting auctions are harder to analyze because of the freedom they give to bidders. Each bidder determines the list of bundles as well as prices on same to announce. In price setting auctions, each bidder is limited to announcing which bundles meet their needs at the announced prices.

In many simple environments price setting and quantity setting auctions can be viewed as being 'dual' to one another. The simplest example is the auction of a single object. The popular English auction is an example of a quantity setting auction. Bidders submit prices in succession, with the object tentatively assigned to the current highest bidder. The auction terminates when no one is prepared to top the current high bid. The 'dual' version to this auction has the auctioneer continuously raising the price. Bidders signal their willingness to buy at the current price by keeping their hands raised. The auction terminates the instant a single bidder remains with their hand raised. In fact this dual version of the English auction is used as a stylized model of the English auction itself for the purposes of analysis (see for example Klemperer 2000).

The discussion of iterative auctions is motivated by this 'duality'. Price setting auctions can be viewed as primal-dual algorithms for the underlying winner determination problem. The reverse will also be true. Primal-dual algorithms for CAP can be given a price setting auction interpretation. Such interpretations are not new. Dantzig (1963), specifically offers an auction interpretation for the decomposition algorithm for linear programming. A more recent example is is Bertsekas (1991), who has proposed a collection of dual based algorithms for the class of linear network optimization problems. These algorithms he dubs auctions algorithms. Auction interpretations of algorithms for optimization problems go back at least as far as Walras (see Chapter 17H of the book by Mas-Collel *et al.* 1995) and all have the same flavor. Dual variables are interpreted as prices and the updates on their value that are executed in these algorithms can be interpreted as a form of myopic best response on the part of bidders.

Is there any reason to prefer an iterative auction to a single round sealed bid auction? Yes, and I list them below. This is followed by a dicussion of their strength.

1. They (might) save bidders from specifying their bids for every possible

combination in advance.

2. Winning bidders need not reveal their valuations. This may be attractive in a repeated situation where a bidder may not want information revealed in one round to be used against them in a subsequent round.

3. A sealed bid auction gives the auctioneer more freedom to manipulate the outcome than an iterative auction.[9]

4. An iterative auction can be adapted to dynamic environments where bidders and objects arrive and depart at different times.

5. In settings where bidders have private information that is relevant to other bidders, iterative auctions (with appropriate feedback) allow that information to be revealed.

6. The cognitive burden on bidders, in terms of choosing a strategy, seems to be lower in (some) iterative auctions than their sealed bid counterpart. See for example Kagel and Levin (2001).

One might be able to incoporate these concerns formally into a model of auction design (eg Larson and Sandholm 2001). In any case, for many of the items of the list there seem to be other ways to resolve the difficulty. Secrecy for bidders and credibility of the auctioneer can be handled by the use of a trusted third party. Private information relevant to other bidders can still emerge in a sealed bid auction by enriching the kinds of bids that bidders may submit, see for example Dasgupta and Maskin (2000). In my opinion only the first and last items on the list have real bit.

5.1 Efficient Iterative Auctions

If the goal is efficiency, then a great deal is known about iterative auctions that accomplish this. Recall the winner determination problem (CAP). If its LP relaxation admits an optimal integral solution, then any primal-dual algorithm for the LP relaxation is a candidate for an iterative auction. The dual variable associated with the constraint

$$\sum_{S \ni i} \sum_{j \in N} y(S, j) \leq 1$$

[9] Implicit in this is that an iterative auction is conducted publically.

can be interpreted as the price of object $i \in M$. The dual variable associated with the constraint

$$\sum_{S \subseteq M} y(S, j) \leq 1$$

can be interpreted as the surplus of agent $j \in N$. In this case, under the assumption that bidders bid sincerely in each iteration, the auction (i.e. primal-dual algorithm) will terminate with the efficient allocation. Examples of this line of inquiry are Kelso and Crawford (1982) and Gul and Stachetti (1999).

What happens if there is an integrality gap in (CAP)? Through the use of auxillary variables one can reformulate (CAP) as a linear program. The dual variables of this extended formulation can interpreted as prices but not on individual objects. Instead they will be prices on subsets of objects and possibly non-annonymous. Examples of extened formulations of (CAP) can be found in Bikhchandani and Ostroy (2000a) and Bikhchandi *et al.* (2001). Parkes and Ungar (2000) develop a primal-dual algorithm for the extended formulation in Bikhchandani and Ostroy (2000a) and give it an auction interpretation. Again, under the assumption of sincere bidding in each iteration, the auction terminates with the efficient allocation.

Is it possible to drop the assumption of sincere bidding? Can one design an iterative auction in which it is an equilibrium for bidders to bid sincerely and still terminate with the efficient allocation? Yes, under certain conditions, or if one is prepared to stretch the intuitive sense of what an iterative auction is.

In Bikhchandani and Ostroy (2000a) it is shown that when $V(K)$ is submodular, there is an LP formulation for (CAP) in which amongst the set of optimal dual solution is one, which gives a surplus to each bidder equal to their marginal product. Equivalently, the dual variables corresponding to prices coincide with the prices that the bidders would pay under a Vickrey auction. Thus the appropriate primal-dual algorithm for this formulation yields an iterative auction in which bidding sincerely is an equilibrium and the efficient allocation is realized (see Bikhchandani *et al.* (2001) for details). Examples of iterative auctions that fit within this framework are Leonard (1983), Demange *et al.* (1986) Ausubel (1997), de Vries *et al.* (2001). When $V(K)$ is not submodular, then no such formulation exists, so submodularity of $V(K)$ is both necessary and sufficient for this approach to work.

In the absence of submodularity of $V(K)$ is an iterative auction that implements the efficient allocation possible? Ausubel (2000) proposes an iterative auction under the assumption that there is no integrality gap in (CAP).

Here I give a high level description of Ausubel's auction for heterogenous objects. [10].

Since there is no integrality gap in (CAP) any primal-dual algorithm for the linear relaxation of (CAP) can be used to determine the efficient allocation (under the assumption of sincere bidding). Call the auction corresponding this primal-dual algorithm the pd-auction. Ausubel's scheme involves $n + 1$ applications of the pd-auction. First run the pd-auction with all bidders. Then, exclude one bidder, and run another pd-auction. The auction that excludes bidder j, say, does not determine the payments to be made by the remaining bidders, but does determine the payment of bidder j. Thus, under the assumption of sincere bidding, Ausubel's scheme determines the efficient allocation associated with $V(N)$ and $V(N \setminus j)$ for all $j \in N$. From the history of price changes in each of the $n + 1$ auctions Aususbel generates enough information to compute the price that each bidders must pay under the Vickrey auction and this is the price they end up paying.

The interesting feature of Ausubel's auction is the use of a collection of auxillary auctions which do not determine the payments and allocations of those participating in it. This is what I mean by stretching the intuitive sense of an iterative auction. This is not the only strech one can conceive of. Parkes and Ungar (2001) have proposed augmenting their origianl iterative auction with fake bidders to prolong the auction so as to generate enough extra information to determine the Vickrey prices.

- The work to date suggests that an ascending auction exists when CAP is computationally feasible. A theorem along these lines would be nice, but the difficulty is in pinning down a 'good' definition of an ascending auction.

- Is there a situation where reporting the valuation function is hard but reporting the set of bundles that maximize net utility is easy?

- If one is willing to accept less than full economic efficiency, perhaps the range of economic environments that admit practical iterative auctions is widened.

- Many practical iterative auctions deal with the computational constraint, by restricting prices to be linear.[11] This invariably leads to the

[10] This description is based on de Vries *et al.* (2001)

[11] The price of a subset of items is the sum of the prices of the items in the subset.

threshod problem (Bykowsky *et al.* (1995)). Consider three bidders, 1, 2 and 3, and two objects $\{x, y\}$. Suppose:

$$v^1(x,y) = 100,\ v^1(x) = v^1(y) = 0,\ v^2(x) = v^2(y) = 75,\ v^2(x,y) = 0,$$

$$v^3(x) = v^3(y) = 40,\ v^3(x,y) = 0.$$

Here $v^i(\cdot)$ represents the value to bidder i of a particular subset. Notice that the bid that i submits on the set S, $b^i(S)$, need not equal $v^i(S)$.

If the bidders bid truthfully, the auctioneer should award x to 2 and y to 3, say, to maximize his revenue. Notice however that bidder 2 say, under the assumption that bidder 3 continues to bid truthfully, has an incentive to shade his bid down on x and y to, say, 65. Notice that bidders 2 and 3 still win but bidder 2 pays less. This argument applies to bidder 3 as well. However, if they both shade their bids downwards they can end up losing the auction. This feature of combinatorial auctions is called the 'threshold problem': a collection of bidders whose combined valuation for distinct portions of a subset of items exceeds the bid submitted on that subset by some other bidder. It may be difficult for them to coordinate their bids to outbid the large bidder on that subset. The basic problem is that the bidders 2 and 3 must decide how to divide $75 + 40 - 100$ between them. Every split can be rationalized as the equilibrium of an appropriate bargaining game. How does one organize the auction to alleviate it?

REFERENCES

L. M. Ausubel, "An Efficient Ascending-Bid Auction for Multiple Objects", manuscript, University of Maryland, 1997.

L. M. Ausubel, "An Efficient Dynamic Auction for Heterogeneous Commodities", manuscript, University of Maryland, 2000.

A. Archer and . Tardos, "Frugal Path Mechanisms," to appear in *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms* , 2002.

D. P. Bertsekas, Linear Network Optimization: Algorithms and Codes, MIT press, Cambridge, MA, 1991.

S. Bikhchandani, S. de Vries, J. Schummer and R. V. Vohra, “Linear Programming and Vickrey Auctions”, to appear in, “Mathematics of the Internet: E-Auction and Markets”, Springer, “The IMA Volumes in Mathematics and its Applications”, Vol. ??, 2001.

S. Bikchandani and J. M. Ostroy, “The Package Assignment Model”, manuscript, 1998.

S. Bikchandani and J. Mamer, “Competitive Equilibrium in an Exchange Economy with Indivisibilities”, *Journal of Economic Theory*, vol. 74, no. 2, 385–413, 1997.

M. M. Bykowsky, R. J. Cull and J. O. Ledyard, “Mutually Destructive Bidding: The FCC Auction Design Problem”, *Journal of Regulatory Economics*, 17, 205–228, 2000.

C. G. Caplice, “An Optimization Based Bidding Process: A New Framework for Shipper-Carrier Relationships”, Thesis in the field of Transportation and Logistics Systems submitted to the Department of Civil and Environmental Engineering, School of Engineering, MIT, June 1996.

W. Conen and T. Sandholm, “Minimal Preference Elicitation in Combinatorial Auctions”, *IJCAI-2001 Workshop on Economic Agents, Models, and Mechanisms*, 2001.

G. B. Dantzig, Linear Programming and Extensions, Princeton University Press, New Jersey, 1963.

A. Davenport and J Kalagnanam, “Price negotiations for procurement of direct inputs”, IBM Technical Report 98774 2001.

G. Demange, D. Gale and M. Sotomayor, “Multi-item Auctions”, *Journal of Political Economy*, 94, 863-872, 1986.

D. Fudenberg and J. Tirole, Game Theory, MIT Press, 1992.

S. de Vries, J. Schummer and R. V. Vohra, “An Ascending Vickrey Auction for Heterogenous Objects”, in preparation, 2001.

F. Gul and E. Stachetti, “English and Double Auctions with Differentiated Commodities”, *Journal of Economic theory*, 92, 66–95, 2000.

J. Hershberger and S. Suri, “Vickrey Prices and Shortest Paths: What is an edge worth?” Proc. of *42nd Annual Symposium on Foundations of Computer Science (FOCS)*, Oct 2001.

P. Jehiel, B. Moldovanu and E. Stacchetti, “Multidimensional Mechanism

Design for Auctions with Externalities", *Journal of Economic Theory* 85(2), 258–294, 1999.

P. Jehiel, B. Moldovanu, "Efficient Design with Inderdependent Valuations", *Econometrica*, vol. 69, 1237-1260, 2001.

J. H. Kagel and D. Levin, "Behavior in Multi-Unit Demand Auctions: Experiments with Uniform Price and Dynamic Vickrey Auctions", *Econometrica*, 69, 413-454, 2001.

V. Krishna and M. Perry, "Efficient Mechanism Design", manuscript, 1998.

K. Larson and T. Sandholm, " Costly Valuation Computation in Auctions: Deliberation Equilibrium", *Theoretical Aspects of Reasoning about Knowledge (TARK)* 169-182, 2001.

D. Lehmann, L. O'Callaghan and Y. Shoham, "Truth Revelation in Rapid, Approximately Efficient Combinatorial Auctions", manuscript, 1999.

K. Leyton-Brown, M. Pearson and Y. Shoham, "Towards a Universal Test suite for Combinatorial Auction Algorithms, *ACM Conference on Electronic Commerce*, 2000a.

A. Mas-Collel, M. D. Whinston and J. R. Green, Microeconomic Theory, Oxford Univ. Press, NY, 1995.

N. Nisan, "Bidding and Allocation in Combinatorial Auctions", *to appear in EC2000*, 2000.

N. Nisan, "Communication Complexity of Combinatorial Auctions", manuscript, 2001.

M. Olson, D. Porter, J. Ledyard, J. Swanson and D. Torma, "The First Use of a Combined Value Auction for Transportation Services", Social Science Working Paper No. 1093, California Institute of Technology, March 2000.

S. J. Rassenti, V. L. Smith and R. L. Bulfin, "A Combinatorial Auction Mechanism for Airport Time Slot Allocation", *Bell Journal of Economics*, vol. 13, No. 2, 402–417, 1982.

J-C. Rochet and L. Stole, "The Economics of Multidimensional Screening", manuscript, 2000.

M. H. Rothkopf, A. Pekec and R. M. Harstad, "Computationally Manageable Combinatorial Auctions", *Management Science*, vol. 44, No. 8, 1131–1147, 1998.

T. Sandholm, S. Suri, A. Gilpin and D. Levine, "CABOB: A Fast Optimal

Algorithm for Combinatorial Auctions", *Proc. 17th IJCAI* 2001.

T. Sandholm and S. Suri, "Side Constraints and Non-Price Attributes in Combinatorial Markets",*IJCAI Workshop on Distributed Constraint Reasoning,* 2001.

J. Schummer, "Manipulation through Bribes," *Journal of Economic Theory* 91, 180-198, 2000.

S. R. Williams, "A Characterization of Efficient, Bayesian Incentive Compatible Mechanisms", *Economic Theory*, vol. 14, 155–180, 1999.

M. Yokoo, "Robust Combinatorial Auction Protocol against False-name Bids," *Artificial Intelligence Journal*, forthcoming.