Elicitation of Honest Preferences for the Assignment of Individuals to Positions

Herman B. Leonard

Harvard University

The problem of eliciting honest preferences from individuals who must be assigned to a set of positions is considered. Individuals know that they will be charged for the positions to which they are assigned. A set of prices that provide no incentive for the individual to misrepresent his preferences is suggested. It is shown that these prices constitute an element of the optimal solution to the dual of a linear programming assignment problem. Both the optimal allocation and the prices to be charged can be derived by solving two linear programming problems once preferences have been elicited. The procedure can usefully be viewed as a simulation of a competitive market under conditions where such a market cannot be expected to function well. It results in an efficient allocation where all resources are valued at their opportunity costs and "consumer surplus" is maximized; its outcome thus has the desirable properties of competitive market equilibria.

In 10 "frontier area" Outer Continental Shelf lease sales conducted by the Department of the Interior between 1970 and 1980 recently reviewed in a study requested by Congress, a total of 1,562 tracts were offered for lease. Only 634 received bids by qualified development companies even though all had been selected for public offering

Thanks are due to a number of colleagues who have helped to make this enjoyable. Richard Zeckhauser and Curt Monash provided both a number of helpful comments and a good deal of encouragement. Ludo Van der Heyden provided a number of useful comments at various stages. John Pratt provided a crucial observation. Nicolaus Tideman, Eric Maskin, and John Riley were particularly helpful in suggesting relevant parts of the literature. An anonymous referee suggested several interesting contextual observations. As always, unfortunately, all of the errors are my own.

[*Journal of Political Economy*, 1983, vol. 91, no. 3]

through industry nomination. Lease Sale Number 43, of the Georgia Embayment area, was conducted on March 28, 1978, in Savannah. Eleven companies participated, and 224 tracts were offered. Only 57 tracts received *any* bids, only 11 tracts received more than two, and no tract received more than six. Even so, 43 tracts were leased and over $100 million changed hands through the sale. While the Georgia Embayment sale is not completely typical, it is fair to surmise that, at least in some instances, disposition of large amounts of public resources is carried out in markets that would have to be characterized as "thin" by virtually any standard.

A great variety of problems considered in economics can be stated in terms of the general "assignment" problem of efficiently matching a set of "individuals" to an equal number of "positions." The marketplace can be thought of as a means of assigning consumers to bundles of goods and services and producers to sets of factor inputs. When each agent is a small part of the market, when large numbers of agents are involved, and when the commodities marketed are homogeneous, the competitive market has many desirable characteristics. But many important assignment problems, like the offshore lease sales just mentioned, involve relatively small numbers of agents, or heterogeneous commodities, or both, so that "markets" used to resolve them would be "thin." Students are matched with dormitory rooms. Employees are assigned to work schedules. Communities are assigned to hazardous waste disposal sites.

The central task of this paper is to develop a workable quasi-market allocation procedure to use in "assignment" problems when the market would otherwise be too thin. If individuals' preferences are known, then the problem of finding the efficient match is not difficult.[1] In general, however, preferences are not known, and the problem is more complicated. We must at the same time *elicit* preferences and *use* them in making assignments. Since the participants will know that their reported preferences affect the assignment, it is not trivial to devise a procedure that requires only information that individuals will voluntarily report. In the language of game theory, we want a procedure that is "incentive compatible," that is, one in which everyone has an incentive to do what is required to guarantee an efficient outcome.

The general problem of incentive-compatible allocation has attracted wide attention and is well understood.[2] Groves (1973) and Clarke (1971) have independently provided very general treatments

[1] For example, see Dantzig (1963). Algorithms for solving the assignment problem have been extensively studied.

[2] See, e.g., Green and Laffont (1979) for a detailed discussion.

of a wide class of problems of which the assignment problem is a special case. The Groves-Clarke (GC) solution involves defining a series of charges, unique up to a set of transfers among individuals; these charges are known to be the only fully incentive-compatible "prices" available.[3] A vital characteristic of the GC solution is that honest announcement of preferences is a dominant strategy for all players. It is therefore "strongly individually incentive compatible."

The very powerful and general observations developed by GC have been applied almost exclusively to the problem of revealing demand for public goods, presumably because private markets cannot efficiently provide public goods, and we therefore need some demand-revealing mechanism to substitute for market provision (see, e.g., Tideman and Tullock 1976; Groves and Ledyard 1977). By contrast, a central concern of economic theory has been to show that competitive markets provide private goods efficiently and that no additional demand-revealing mechanism is required. The theory shows that the competitive marketplace is an (almost perfectly) incentive-compatible demand-revealing allocation mechanism, whose (nearly full) incentive compatibility comes from the fact that the scale of each participant is small relative to the market as a whole, so that no individual can affect market prices. When that assumption is not met, we need to look for an alternative incentive-compatible allocation mechanism. The GC procedure as applied to public goods in effect "simulates" a competitive market where one cannot exist because benefits of consumption are not individual. In this work I develop a similar "simulation" of a competitive market where one cannot be assumed to exist because there are too few interested participants. The problem is posed schematically in the form of the "assignment problem" of matching individuals to positions. Since only one individual can be assigned to any position, I am assuming that the goods being allocated are "private" goods, whose benefits are individually appropriable.

Because the assignment problem is a special case of the general problem treated by Groves and Clarke, I could proceed by simply applying the GC procedure in this context. This would leave the usual computational problem of finding GC prices once preferences are elicited. Instead I shall find incentive-compatible charges for the assignment problem from first principles and then demonstrate that the resulting charges are in fact those suggested by GC. This approach

[3] Jacob Marschak was reportedly the first to make the elegant observation that incentive compatibility requires that the price charged to an individual not depend on his statements of preference. Vickrey (1961) is apparently the first to have developed this idea in a published paper. Groves (1973) and Clarke (1971) provided very general treatments.

has two advantages. First, it provides a simple and direct computational scheme for deriving the appropriate charges. Second, it shows the link between the GC scheme and a thoroughly understood mechanism (linear programming) and may thus provide an additional basis for understanding the GC suggestion.

The basic idea of the solution is directly analogous to the results of Groves and Clarke in their more general context. I show that of all the prices in the dual solution to the assignment linear programming problem—that is, all the prices that support the efficient allocation of individuals to positions—only one set is incentive compatible. This is the set obtained by pricing each position on the basis of valuations stated only by those *not* assigned to it. The value assigned to each position is the difference between the collective valuation (of all individuals not assigned to it) of *all* the positions and their valuation of all *but* the one in question. Because this valuation takes no account of the individual efficiently assigned to the position—the "highest bidder" of sorts—it results in the *lowest* set of possible valuations that sustain the optimal allocation.

Hylland and Zeckhauser (1979) discuss a variant of the assignment problem where there is no external "medium of exchange" that links the problem at hand to the "outside world." As a consequence of this constraint, their procedure is not completely incentive compatible. Instead, they argue that where "large numbers" of individuals are involved, so that no individual is very influential, their procedure is essentially compatible with honest preference revelation—that is, the incentive to misrepresent becomes weak.

In this work, I explicitly assume that there is a "currency" in which individuals can be "charged" for their assigned positions. Introducing this currency makes it possible to devise a procedure which is fully incentive compatible, even with only a small number of participants. The procedure can be thought of as a way to simulate a working competitive market when some of the usual reasons for presuming that the market will function well do not apply.

Like Hylland and Zeckhauser, I presume the existence of a benevolent auctioneer with a markedly different objective than that typically assumed in the modern auctions literature. In most recent work, the problem has been to maximize the return to the auctioneer himself. Depending on the form of the information available to him, this objective could sometimes be served by not selling some of the items. In the treatment given here, the auctioneer starts with no knowledge of preferences and makes no attempt to affect the total payment received from selling the positions. His sole objective is to ensure that the allocation is efficient; the revenues received are treated as irrelevant. An interesting implication of this combination of reliance on elicited preferences and desire for efficiency is that, far from being maximized, the returns to the "auctioneer" in this work turn out to be the *minimum* possible returns compatible with the optimal allocation. Such benevolence on the part of an auctioneer can perhaps be expected realistically only when this role is played by the "government."

I. The Assignment Problem

A set of risk-neutral individuals must be matched with an equal number of positions, one to each position. Some positions may be null, that is, valued by no one, and others may be identical. I wish to allocate the positions efficiently, by which I mean that participants should not wish to arrange trades among themselves after the allocation is made. I assume that there is a currency which has value outside the specific assignment problem under consideration. This currency can be used to charge individuals for positions to which they are assigned. The problem is to devise a procedure that (1) elicits honest preferences, (2) results in an efficient allocation of individuals to positions, and (3) charges prices that cause every individual to choose his best position.

I begin by assuming that I have elicited honest preferences and proceed to solve the assignment problem. I will later ensure that the solution is compatible with honest preference revelation. The assignment portion of this "elicitation and matching" problem was originally stated formally as a linear programming problem by Koopmans and Beckmann (1957); an alternative game-theoretic formulation was considered by Shapley and Shubik (1972). Let i index individuals and j index positions for i and $j = 1–N$, and let $h_{ij} \geqslant 0$ denote the valuation of position j by individual i. Then the problem is to choose x_{ij} so as to

$$\text{maximize} \sum_{ij} x_{ij} h_{ij} \tag{1}$$

subject to

$$a) \sum_{i} x_{ij} \leqslant 1 \text{ for any } j,$$

$$b) \sum_{j} x_{ij} \leqslant 1 \text{ for any } i,$$

$$c) \; x_{ij} \geqslant 0 \text{ for any } i, j.$$

Koopmans and Beckmann (KB) observe that there is a solution to this problem which involves only values of zero and one for the activity variables x_{ij}. These variables denote whether individual i is assigned to position j. The first constraint states that only one individual can occupy each position; the second constraint states that each individual can occupy only one position.[4] Since there is a finite number of possible assignments, there must be an optimal feasible solution. Given the preferences h_{ij}, the solution can easily be found using the simplex method or any of a number of algorithms that take account of the special structure of the "assignment" linear programming problem (see, e.g., Hillier and Lieberman 1974).

Having solved the assignment problem (conditional on knowing the true preferences), KB turned to the problem of finding "prices" that "support" the optimal allocation. A well-known property of linear programming is that the "dual" program provides market-clearing prices for the "resources" used in the primal problem. In this case, the resources are the individuals and the positions available. Thus, the dual program provides "shadow prices" for both the individuals and the positions. We might imagine charging for assignments to positions in the form of (1) an "admissions fee," s_i, specified for each individual, charged for participating in the allocation procedure, and (2) a "positions" fee, v_j, specified for each position, charged to whoever is assigned to it. The charges can be derived from the dual problem:

$$\text{Minimize} \sum_i s_i + \sum_j v_j \tag{2}$$

subject to

$$a)\ s_i + v_j \geqslant h_{ij} \text{ for any } i, j,$$

$$b)\ s_i, v_j \geqslant 0 \text{ for any } i, j.$$

Intuitively, the dual problem says that the resources are to be assigned the minimum values compatible with the fact that they are valued by the recipients. Indeed, as KB point out, the resources are assigned prices related to their values in their best possible uses. This "best use" valuation is a characteristic of any competitive market equilibrium.

Because we know that the primal problem has a solution, we know also that the dual must have an optimal solution, and it must attain the same value. KB note that, due to the multiple degeneracy of the primal assignment problem, there will typically be multiple solutions

[4] These constraints are a consequence of the way I have defined "individuals" and "positions." It should be noted that nothing in this formulation requires that all the positions differ from each other. The problem is not very hard or very interesting if all of them are the same, but some of them may be identical.

TABLE 1

ASSIGNMENT OF TWO INDIVIDUALS

	Individual 1	Individual 2
Position 1	12	11√
Position 2	7√	4

to the dual problem; *any* set of these dual solutions provide "market-clearing prices"—if we announce a set of these prices, every individual will be satisfied "buying" the position to which he has been optimally assigned.

It is important to note that the chosen allocation does not depend on the "personal charges" s_i, provided that they do not vary with the position chosen. The shadow price s_i is like an individual admission fee; it could be used as a means of extracting the "surplus" that individual i would otherwise receive. Whatever the absolute level of s_i, individual i will be made at least as well off by choosing the position to which he is optimally assigned as by choosing any other position. Thus, the shadow prices on individuals do not enter into the allocation decision. Intuitively, this is because individuals choose positions and not vice versa. The implication is that the only prices we need to set correctly are the position shadow prices v_j; the s_i's are irrelevant and typically will all be set to zero.

The multiplicity of price systems that support an optimal assignment can be illustrated in a simple example. Table 1 shows an assignment of two individuals to positions. The table shows the true valuation of each individual for each position; the checks designate the optimal assignment. This assignment is supported by any fixed set of individual charges s_i and any position prices for which

$$12 - v_1 \leq 7 - v_2 \tag{3a}$$

and

$$11 - v_1 \geq 4 - v_2 \tag{3b}$$

or, equivalently,

$$7 \geq v_1 - v_2 \geq 5. \tag{4}$$

In addition, we must have $v_i \geq 0$, and since the individual charges $h_{ij} - v_j$ must be at least zero, we also require that $v_1 \leq 11$ and $v_2 \leq 7$. The shaded area in figure 1 illustrates the region of price combinations (v_1, v_2) that meet these conditions.

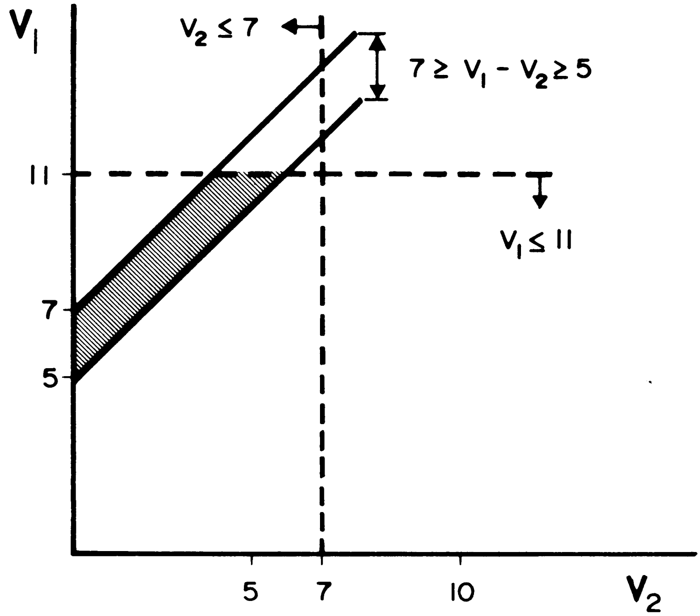

FIG. 1.—Market-clearing price combinations

KB were unconcerned by the multiplicity of prices they found that would support the efficient assignment.[5] Moreover, because they were not considering the process through which prices were established, they did not inquire into whether individuals could be expected to reveal their true valuations if they knew how prices were to be determined. As the next section will demonstrate, addressing one of these concerns addresses the other as well: adding the requirement that the prices announced be incentive compatible leads to a unique and intuitively sensible choice from among the available market-clearing price systems.

II. Incentive Compatibility

Given a region of price combinations that support the optimal allocation—presuming elicitation of honest preferences—I shall now turn to the incentive compatibility problem to identify a set of these prices which not only "clear the market" but also provide an incentive for

[5] Shapley and Shubik (1972) consider the geometry of the set of solutions to the dual of a related game-theoretic problem in some detail. Their central concern is to show that the set of dual solutions is coincident with the core of the game they describe. They do not consider the issue of price determination or incentive compatibility.

individuals to reveal their true valuations. Since the "individual charges" s_i do not affect the allocation, I shall concentrate on finding a set of incentive-compatible "positions prices" v_j.

Since we know that shadow prices for positions from the dual program are market-clearing prices, we begin by examining a set of them to see whether they might be incentive compatible. Intuitively, a shadow price gives the change in the value of the objective function if the constraint it represents is loosened or tightened. In the assignment problem, as in any linear programming problem, the value lost when a constraint is tightened may not be the same as that gained when it is relaxed. Consider the shadow price of a position generated by tightening the constraint—that is, by eliminating the position. Denote by V_I^P the value of the optimal assignment of the individuals in set I to positions in set P. Then the shadow price generated by tightening the constraint on position j is given by the difference between the value of all the positions to all the individuals, V_I^P, and the value of positions other than j to all the individuals, V_I^{P-j}:

$$\bar{v}_j = V_I^P - V_I^{P-j}. \tag{5}$$

If we announced that we were going to charge $\bar{v}_j$, as defined above, for position j, we could not reasonably expect individuals to respond with their true valuations, because the charge depends on the stated valuations of *every* individual, and in particular on the stated valuation of the individual who pays it. Consider individual i, who is optimally assigned to position j. His "surplus value" after paying the charge $\bar{v}_j$ will be

$$h_{ij} - (V_I^P - V_I^{P-j}). \tag{6}$$

Since both V_I^P and V_I^{P-j} depend on i's stated valuations, he has a potential incentive to misrepresent his valuations. Thus, while prices $\bar{v}_j$ defined as above *do* support the optimal allocation if preferences are reported honestly, they provide an incentive for individuals not to represent their preferences truthfully.

Suppose instead we focus on the shadow price generated by *relaxing* the constraint—that is, that corresponding to *adding* another position of type j. This shadow price is the difference between the value to all individuals of all positions plus another position of type j, V_I^{P+j}, and the value of the current positions, V_I^P:

$$\underline{v}_j = V_I^{P+j} - V_I^P. \tag{7}$$

Again, this appears to depend on the valuations stated by every individual. In fact, however, it does not.

To see this, first note that if the optimal allocation of the individuals I to positions P calls for individual i to be allocated to position j, then if

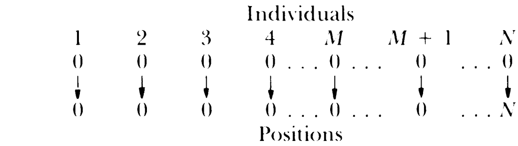

A. The optimal allocation before the new position is added

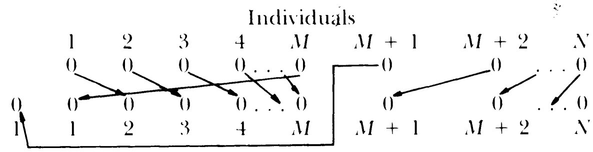

B. A reallocation with one new position and with individual 1 assigned to neither the old nor the new position of type 1. The shuffle among the first M individuals and positions was available before the new position was added and was not utilized as part of the optimal assignment. Hence it cannot be part of the optimal assignment now. This cannot, therefore, constitute the optimal reassignment.

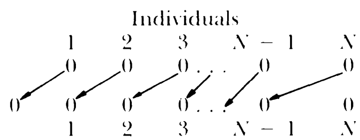

C. An optimal reassignment with a new position

FIG. 2

another position of type j is added individual i will continue to occupy a position of type j. Providing a new position of type j cannot be better for individual i unless someone else moves into the new position j and i moves up to some position that becomes vacant. But this shuffle was available before the new position was added. It was not optimal to reshuffle before; therefore, it cannot be optimal now. Thus, individual i must remain in position j. Figure 2 illustrates this schematically.[6]

As V_I^{P+j} must involve assigning individual i to position j, which has value h_{ij},

$$V_I^{P+j} = V_{I-i}^{P+j-j} + h_{ij} = V_{I-i}^{P} + h_{ij}. \tag{8}$$

Similarly, V_I^P involves assigning i to j, and so

$$V_I^P = V_{I-i}^{P-j} + h_{ij}. \tag{9}$$

Therefore,

$$\underline{v}_j = V_I^{P+j} - V_I^P = (V_{I-i}^{P} + h_{ij}) - (V_{I-i}^{P-j} + h_{ij}), \tag{10}$$

[6] I am indebted to John Pratt for suggesting this compact argument.

so that

$$\underline{v}_j = V^P_{I-i} - V^{P-j}_{I-i}. \tag{11}$$

This does not depend on the values that individual i stated. That is, $\underline{v}_j$ is the difference between the optimal valuation achievable by allocating everyone *except* person i to all the positions and the optimal valuation obtainable by allocating everyone except person i to all positions other than j. If I announce my intention to charge $\underline{v}_j$ for position j, individual i's surplus will be

$$h_{ij} - (V^P_{I-i} - V^{P-j}_{I-i}), \tag{12}$$

which does not depend on the valuation he *states* for the position he occupies; rather, it depends on his true valuation h_{ij} and on statements of value made by others. Thus, I have provided no incentive for him to misrepresent his valuation.

These prices also provide a positive incentive for him to tell the truth. If individual i is optimally assigned to position j and truthfully states his valuations, he will pay $\underline{v}_j = V^P_{I-i} - V^{P-j}_{I-i}$ for it and receive surplus

$$\bar{s}_i = h_{ij} - (V^P_{I-i} - V^{P-j}_{I-i}). \tag{13}$$

If he misrepresents his valuations and does not affect the assignment, his surplus will be the same and he will have gained nothing. Suppose he *does* affect the assignment by misrepresenting his valuations and is assigned instead to position k. He will now receive h_{ik} and will pay $V^P_{I-i} - V^{P-k}_{I-i}$. To see that he cannot gain from such a misrepresentation, we need to show that his new surplus is at most what he receives if he tells the truth, or that

$$h_{ij} - (V^P_{I-i} - V^{P-j}_{I-i}) \geq h_{ik} - (V^P_{I-i} - V^{P-k}_{I-i}). \tag{14}$$

This is equivalent to showing that

$$h_{ij} + V^{P-j}_{I-i} \geq h_{ik} + V^{P-k}_{I-i}. \tag{15}$$

The value on the left is simply V^P_I, the highest value attainable by assignment of individuals I to positions P. The value on the right is the value of the best arrangement of all individuals other than i to positions other than k plus the value of assigning i to k. Because i is not optimally assigned to k, the value of the *constrained* assignment on the right, with individual i in position k, can at best be as large as the unconstrained assignment on the left, with i optimally located in position j. Thus, (15) and (14) both hold; the surplus from telling the truth is always at least as great as that from lying and typically is greater.

These prices not only provide no positive incentive to misrepresent, they also provide a positive incentive to tell the truth; they are fully incentive compatible.[7] As the GC prices give the unique fully incentive-compatible prices, these prices must be the GC prices.[8] Using the shadow prices represented by $\underline{v}_j$ provides an incentive-compatible set of market-clearing prices.

III. Computation of Incentive-compatible Prices

The preceding section identified a set of incentive-compatible prices based on the dual to the assignment problem and showed that these prices must be the GC prices. It is straightforward but tedious to show that these "position" prices, together with a similarly defined set of "individual" charges, constitute one optimal solution to the dual problem; this is demonstrated in an appendix available from the author. What remains is to develop a computational procedure for selecting the unique set of incentive-compatible shadow prices from the class of all dual solutions.

We could find these prices by solving for the valuations that constitute their definitions. For example, to find $\underline{v}_j$ we would need to know V^P_{I-i} and V^{P-j}_{I-i}. The latter value can be computed easily from V^P_I (by subtracting h_{ij}), and V^P_I must be computed in solving for the optimal assignment. V^P_{I-i} can be obtained from the solution of the assignment of individuals $I - i$ to positions P. Thus, to find the prices directly, we would need to solve the primal problem of allocating the individuals I to positions P, and the N problems of allocating individuals $I - i$ to P, a total of $N + 1$ linear programming problems. It would be nice to find a more efficient computational procedure.

An alternative would be to identify from first principles which element of the set of solutions to the dual constitutes the incentive-compatible prices. The fact that the prices $\underline{v}_j$ depend only on the valuations stated by $N - 1$ individuals, while the prices $\bar{v}_j$ depend on the statements of all N individuals, suggests that the prices we seek are lower than other solutions to the dual. Indeed, we might suspect that the relevant prices may be the market-clearing positions prices with the lowest sum. In effect, the prices identified for each position represent the outcome of competition only among individuals not assigned to that position; because they were not assigned to it, their valuation must be lower than that of the recipient. The formal demonstration is tedious, but the GC prices are in fact the lowest positions prices that

[7] To be completely precise, I should perhaps note that the payoff $h_{ij} - \underline{v}_j$ is *locally* flat with respect to announcements of h_{ij}; as long as it does not affect the allocation, it will not hurt an individual to misrepresent his valuation.

[8] We could also simply note that the definition of $\underline{v}_j$ is the GC definition.

clear the market—that is, there is no set of positions prices having a lower sum that support the optimal allocation; the appendices (available from the author) present a formal argument.

Because the condition that the prices we seek are the "lowest" prices that solve the dual problem is a *linear* constraint, we can find the prices we want by solving two sequential linear programming problems. First, using any efficient assignment algorithm, we solve the primal assignment problem:

$$\text{Maximize} \sum_{ij} x_{ij}h_{ij} \tag{16}$$

subject to

$$a)\ \sum_{i} x_{ij} \leqslant 1,$$

$$b)\ \sum_{j} x_{ij} \leqslant 1,$$

$$c)\ x_{ij} \geqslant 0.$$

This provides both the optimal assignment and the value of that assignment, V_I^P. I now wish to find that set of prices $\underline{v}_j$ that minimizes $\Sigma_j v_j$ in the class of prices which meet the other conditions of the dual problem and which attain this optimum value. This second problem is to

$$\text{minimize} \sum_{j} v_j \tag{17}$$

subject to

$$a)\ s_i + v_j \geqslant h_{ij},$$

$$b)\ s_i,\ v_j \geqslant 0,$$

$$c)\ \sum_{i} s_i + \sum_{j} v_j = V_I^P.$$

The first and second conditions state that the solution $\underline{v}_j$ must still solve the dual problem; the third condition states that the solution must attain the optimum valuation.[9] Thus, we can solve for the optimum assignment and the incentive-compatible prices that support it by solving only two linear programming problems.

It is worth noting that minimizing $\Sigma_j v_j$ while holding $\Sigma_i s_i + \Sigma_j v_j$ constant maximizes $\Sigma_i s_i$. The s_i are consumer surpluses—the differ-

[9] I am indebted to Ludo Van der Heyden for suggesting this formulation.

ence between the values received (h_{ij}) and the prices charged for positions (v_j). Thus, using the GC prices results in leaving in the hands of participants the maximum possible consumer surplus.

IV. Interpretation and Illustrations of the Suggested Procedure

We now have an incentive-compatible procedure that simultaneously elicits honest valuations and uses the valuations to find an efficient allocation of individuals to positions together with a relatively efficient computational algorithm for finding the required incentive-compatible charges; we turn now to interpreting and illustrating the suggested method.

The most straightforward interpretation of the procedure is that it simulates the functioning of a competitive market allocation process in thin markets not likely to yield competitive outcomes; in this interpretation, we might refer to the mechanism as a "pseudomarket" procedure.

The procedure is analogous to a competitive market in several ways. First, it is fully incentive compatible, a characteristic it shares only with traditional markets in which each agent is so small a part as to be a true price taker. Second, its resulting allocation is fully efficient. Third, it provides valuations based on best use—all resources are valued at their opportunity cost, the value they would generate in the next alternative use. Finally, the prices charged for commodities are the market-clearing prices with the lowest sum—charging these prices leaves in the hands of consumers the maximum possible surplus given that they must be persuaded to purchase their optimally assigned positions.

A second interpretation of this procedure can be derived from its application to a special case of the assignment problem where only one "position" is of value. Consider a sale of a single Rembrandt painting. Of various ways available to carry out this "assignment," the well-known "second price" auction first described by Vickrey (1961) has the nice property of incentive compatibility. The second price auction is conducted by soliciting sealed bids for the painting and "selling" it to the highest bidder at the price offered by the second highest bidder. It is a simple matter to show that the procedure developed here gives precisely the same result—our procedure reduces exactly to the second price auction mechanism when only one position is of value to everyone.

We can interpret this procedure, then, as a generalized form of the second price auction where a variety of goods (positions) are to be simultaneously auctioned. It retains the incentive compatibility of the single-commodity second price auction, while embodying the fact that individuals may have *exclusive* demands. For example, A may want either a silver tray or a gold necklace, but not both. To represent this, A would enter the allocation procedure as one individual, who could be assigned at most to one position.

It should be noted that the formal assignment problem treated here is quite general. In the usual conception of an "assignment," all individuals and all positions are different. We have assumed that each "individual" receives only one item and that each item goes to only one individual. But nothing in these assumptions dictates that two or more items could not be identical or that participants could not enter the procedure under several aliases. If, in the example above, B is willing to buy either the silver tray or the gold necklace separately, he can enter the procedure as two individuals, B1 and B2, each stating an honest preference for only one of the items. Naturally, he will collude in his B1 bid for item 1 with himself (as B2) by stating a zero price for the second item, and vice versa. But this is exactly what we want him to do, for it implies that the price he pays for either item (or both) if assigned to him will be determined only by true "other" bidders. The system maintains its incentive compatibility even when some bidders enter more than once or some positions are the same as others.

The assignment problem as treated here and the algorithm developed to resolve it in an incentive-compatible way thus encompass a very broad class of "matching" problems. The only requirement is that the preferences stated by each individual for each item (h_{ij}) must not depend on the outcome of the assignment. Each individual's preferences for the items must be either completely *independent* (my desire to have X is unrelated to whether I have Y) or *exclusive* (I value either X or Y but not both). Preferences in which one item is more valuable in the presence of another—as, for example, where one needs two plates to complete a set—cannot be represented. This means that the procedure cannot be used for multiple identical items for which individuals have downward-sloping demand curves, because their valuations of one unit will depend on how many other units have been assigned to them. The procedure is thus perhaps most useful where a series of items is being "sold" to a series of "bidders" and each bidder is interested in at most one item of each type and values items of different types independently.

When all individuals state independent demands for all goods, then the multiple auction procedure will reduce to a series of independent second price auctions. Because individuals' demands for each item do not depend on what happened to the others, the items can be auctioned in any order. Sequential auctions thus implicitly assume that individuals' demands are *not* exclusive. When some individuals *do* have exclusive demands, independent second price auctions may not elicit honest bids and may thus not assign individuals efficiently. For example, suppose a collector wants to buy either a Titian or a Rembrandt. If the two paintings are auctioned at the same time, say, by "sealed second price bid," he will not honestly represent his valuations because they will be treated as independent when in fact they are not. He will tend to underrepresent his valuation for at least one of the items, for fear of actually receiving both.

In order to ensure an efficient allocation, the mechanism used must be able to embody exclusive demands if individuals are likely to have them. The procedure suggested here can be interpreted as a natural generalization of the single-commodity second price auction mechanism to a setting where multiple goods are being auctioned and exclusive demands are anticipated. In this application, the procedure might be called a "multiple-commodity second price auction."

There are many applications in which such a pseudomarket or multiple-commodity second price auction mechanism might provide a useful alternative to negotiated settlements or an unconstrained market. We briefly sketch three examples in which the procedure suggested here might play a constructive role. The first example is a strict application of the framework we have developed; the other two require interpretive generalizations of the approach.

Example 1. Matching students to dormitory rooms. This is a classical matching problem in which each individual is unique and must be assigned to one and only one position. Assignments of this type are made on a continuing basis. Participants might be allocated points to be used in a series of such problems; seniors can be given more points than freshmen, or students can be left to allocate their points across years as they see fit.

Example 2. Assigning hazardous waste disposal, correctional facility, or other state or federal government sites to communities. This is a generalized matching problem in which bidders can be assigned to more than one position—that is, can enter the procedure as more than one bidder—but probably have bids that are exclusive within a particular type of facility. In this case, of course, most communities will be bidding most for the "null" positions of not having these sites located in their midst. The distribution of points could depend on the socioeconomic status of residents. For example, they might be distributed as an increasing function of the inverse of average income or in proportion to population.

Example 3. Auctions of oil lease tracts for exploration. This is the example with which we began. This is a generalized matching problem in the sense that bidders may wish to be assigned to more than one tract. It is not a strict application of the method outlined here, because the demand by each bidder for every tract might depend on whether he has been assigned to other tracts; many small bidders may not be able to afford more than a few tracts. A fully general bidding arrangement would be to have each bidder state valuations for each tract, together with a budget constraint stating how much they could spend in total. Such bids could then be processed through a (more complicated) GC procedure. The procedure suggested here cannot, of course, fully represent this more general set of constraints. It would, however, constitute a substantial improvement over the current bidding procedure, which assumes that all bids by any company are independent. Many small oil companies cannot reasonably expect to pursue exploration in more than a few tracts at a time and thus may hesitate to bid on more than a few at any one auction. If they could state that some of their bids are exclusive of one another, the bidding for all tracts would become more competitive. While it does not provide a fully general solution in this instance, the procedure suggested here dominates that currently used, for it includes fully independent bidding as a special case.

In some instances, the appropriate currency for stating valuations and for collecting charges will be money; in other cases, money may be regarded as a poor choice of bases. In situations in which the intent is not actually to sell the positions, the fact that the allocation procedure collects revenue may be an embarrassment. It is not difficult to devise methods by which "most," though not all, of the revenue might be rebated in a way that does not destroy the incentive compatibility of the procedure.[10] Whatever currency is adopted must not derive its value exclusively from any particular assignment problem; it must have a "shadow price" or "opportunity cost" from uses *outside* a given problem in order to serve as a comprehensive basis for stating valuations within it.

[10] The "revenue" raised by this allocation scheme may be an embarrassment in problems in which the sole purpose is determining the appropriate assignment. We cannot rebate all of the revenue raised, for the total revenue depends on the statements of every individual. We can, however, rebate to each individual a share of the revenue that would have been raised if he had not been in the problem. This does not depend on his stated valuations, so knowing that he is going to receive such a rebate will not alter his incentive to reveal his preferences. The total revenue that would have been raised if individual i had not been in the process would be $R_i = \Sigma_j(V^P_{I-i-l} - V^{P-j}_{I-i-l})$, where l is the individual optimally assigned to position j in the absence of individual i. Since the actual allocation *did* include consideration of i's valuation, it must result in charges at least as high as those that would have been charged in i's absence. If we choose to rebate to each individual a share α_i of the rebatable revenues, where $\Sigma_i\alpha_i \leq 1$, we can give individual i a rebate of $\alpha_i R_i$; this will always leave some funds as net collections by the system (since each R_i is at most as large as the actual revenue raised), but it does succeed in rebating some of the revenue in an incentive-compatible way.

V. Concluding Remarks

Assignment problems are frequently encountered. I have proposed a procedure that elicits honest preferences and uses them both to assign individuals to positions efficiently and to determine prices to be charged for the positions. The optimal allocation and the prices to be charged are solutions to a linear programming formulation of the assignment problem given these honestly revealed preferences; we developed an efficient two-step algorithm for conducting the allocation procedure and deriving the appropriate prices.

The mechanism can be interpreted usefully as a simulation of a competitive market allocation process or as a generalization of single-commodity second price auctions to the case of multiple "commodities" for which individuals wish to state either completely independent or completely exclusive demands. An obvious application is to situations in which the "assignments" are actually sales of commodities; the system can be no worse and will in general be better than "independent" second price auctions. The system can also be applied to many assignment problems in which "auctioning" or "selling" positions might be inappropriate. In these cases, some currency other than money can be used to state valuations and to collect charges for positions. The only requirement is that the currency have a value not determined solely in the context of the given assignment problem.

The procedure described here has the distinct disadvantage that it relies for "market discipline" on an external medium of exchange. In a society in which one such medium, money, is increasingly unpopular as a basis for allocation for some commodities and positions, this may be a serious shortcoming. The incentive compatibility of the procedure was purchased at the expense of introducing this medium of exchange. Hylland and Zeckhauser (1979) have proposed a procedure with no currency that is essentially incentive compatible when the number of participants is "large." The mechanism proposed here may thus be most useful in contexts where the number of participants is "small" and the market therefore "thin."

References

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

Dantzig, George B. *Linear Programming and Extensions.* Princeton, N.J.: Princeton Univ. Press, 1963.

Green, Jerry R., and Laffont, Jean-Jacques. *Incentives in Public Decision-Making.* Amsterdam: North-Holland, 1979.

Groves, Theodore. "Incentives in Teams." *Econometrica* 41 (July 1973): 617–31.

Groves, Theodore, and Ledyard, John O. "Optimal Allocation of Public Goods: A Solution to the 'Free Rider' Problem." *Econometrica* 45 (May 1977): 783–809.

Hillier, Frederick S., and Lieberman, Gerald J. *Introduction to Operations Research.* San Francisco: Holden-Day, 1974.

Hylland, Aanund, and Zeckhauser, Richard. "The Efficient Allocation of Individuals to Positions." *J.P.E.* 87 (April 1979): 293–314.

Koopmans, Tjalling C., and Beckmann, Martin J. "Assignment Problems and the Location of Economic Activities." *Econometrica* 25 (January 1957): 53–76.

Shapley, Lloyd S., and Shubik, Martin. "The Assignment Game I: The Core." *Internat. J. Game Theory* 1, no. 2 (1972): 111–30.

Tideman, T. Nicolaus, and Tullock, Gordon. "A New and Superior Process for Making Social Choices." *J.P.E.* 84 (December 1976): 1145–59.

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