

The Efficient Allocation of Individuals to Positions

Author(s): Aanund Hylland and Richard Zeckhauser

Source: *Journal of Political Economy*, Vol. 87, No. 2 (Apr., 1979), pp. 293-314

Published by: The University of Chicago Press

Stable URL: <https://www.jstor.org/stable/1832088>

Accessed: 27-05-2020 17:49 UTC

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

The University of Chicago Press is collaborating with JSTOR to digitize, preserve and extend access to *Journal of Political Economy*

The Efficient Allocation of Individuals to Positions

Aanund Hylland and Richard Zeckhauser

Harvard University

In a variety of contexts, individuals must be allocated to positions with limited capacities. Legislators must be assigned to committees, college students to dormitories, and urban homesteaders to dwellings. (A general class of fair division problems would have the positions represent goods.) This paper examines the general problem of achieving efficient allocations when individuals' preferences are unknown and where (as with a growing number of nonmarket allocation schemes) there is no facilitating external medium of exchange such as money. An implicit market procedure is developed that elicits honest preferences, that assigns individuals efficiently, and that is adaptable to a variety of distributional objectives.

In a variety of contexts, individuals must be allocated to positions with limited capacities. Faculty members or legislators must be assigned to committees, family members must be given household chores, secondary school students must be placed in special-interest high schools, and college students must be assigned to popular courses with limited enrollments. These individuals are indivisible; fractions of them cannot be assigned to different positions.

Ideally, these assignments should be efficient (Pareto optimal). If individual preferences are known, a mechanical application of the

We thank Kenneth Arrow, Robert Klitgaard, Albert Nichols, and Thomas Schelling, participants in seminars at Harvard, Northwestern, and the University of Pennsylvania; a referee; and most particularly, Howard Raiffa for helpful comments. Earlier versions of this paper were circulated as RIAS working paper no. 7, and Kennedy School Discussion Paper 51D, Harvard University. Research support from the Norwegian Research Council for Science and the Humanities (Hylland) and NSF grants SOC77-16602 and SER76-17502 is gratefully acknowledged.

[Journal of Political Economy, 1979, vol. 87, no. 2]

©1979 by The University of Chicago. 0022-3808/79/8702-0003\$02.00

“job-assignment algorithm” will produce an efficient outcome.¹ A central aspect of the problem, however, is that preferences are unknown and can only be discovered by asking the individuals or letting them respond to an assignment procedure. It is then possible that they will respond strategically rather than truthfully. Money (or some other medium of exchange) is the traditional prerequisite for efficient allocation under such conditions. But in the examples given, and in many other important cases, money is not an acceptable instrument.²

This paper examines the general problem of achieving efficient allocations in such circumstances. We employ the paradigm of assigning individuals to jobs, although most applications would involve the assignment of other kinds of entities to other kinds of positions. For example, an important class of fair-division problems can be represented by the model; here our “jobs” would correspond to units of different indivisible goods, with individuals receiving multiple assignments. The capacity constraints on the jobs can be maxima—at most 10 Representatives can serve on the Congressional Energy Committee; or minima—at least 80 white ninth graders must attend West Side High School; or the constraints may have to be met with equality—the appointment committee must have exactly six members and the budget committee exactly four.

This problem, in essence, is to find a social choice mechanism. To achieve satisfactory outcomes, the assignment procedure, that is, the social choice mechanism, must simultaneously (1) elicit honest indications of preferences from the individuals being assigned;³ (2) in light of those preferences, efficiently allocate individuals to positions; (3) meet prescribed distributional objectives. In some cases, it may be desired to treat everybody equally; in others, it may be decided that certain individuals should be systematically favored. The system must be able to handle either situation.

¹ See, for example, Dantzig 1963. To use the job-assignment algorithm, we must also be able to compare and weight different individuals' preferences.

² This is possibly a controversial statement. Some will perhaps claim that it is perfectly acceptable—and indeed in everybody's interest—for example, that well-to-do students buy their way into preferred dormitories. Society's decision not to allow money to be used in many cases of the type we discuss is not the subject of our study and is taken for granted.

³ In a sense, this requirement is redundant. If a mechanism always produces allocations which are efficient relative to people's true preferences, it should not concern us that some individuals have not reported the truth. More generally, one can consider mechanisms in which people respond by choosing a message from a prescribed but arbitrary message space; then the issue of truthful response becomes irrelevant.

We want to construct a mechanism in which each individual has a “canonical” response which depends only on his own preferences (and the rules of the game). That is, we are not satisfied if everybody's response has to depend on everybody's preferences in order for efficiency to be achieved and strategic opportunities to be avoided. When this is the goal, there is no real loss of generality in assuming that the “message space” is the set of possible preferences and the canonical response is telling the truth.

Experience with related problems in social choice theory might suggest that no such procedure exists. Initially, the literature in this area, which grew out of Kenneth Arrow's celebrated Impossibility Theorem, did not consider the question of eliciting honest revelation of preferences; several recent contributions, however, have argued that it is not in general possible to construct nonmanipulable social choice procedures without giving up other desirable properties, such as efficiency.⁴

Our work, on the other hand, has yielded a positive result. In general, if the number of individuals to be assigned is large relative to the capacities of the positions, a procedure is available that elicits honesty, assigns efficiently, and can be adapted to different concepts of equity. Not surprisingly, the mechanism relies on an implicit market.⁵

We are interested in more than the mere existence of a successful procedure. We shall develop one, and explain and illustrate its use in practice.⁶

⁴ See Zeckhauser 1973 and Gibbard 1978.

⁵ There has recently been a growing literature on constructing tax systems which elicit honest preferences for and provide optimal amounts of public goods. This problem is related to the one we are studying as far as the core objectives are concerned. The goal is to construct a procedure which elicits honest responses and deals with them in a manner that produces efficient outcomes. The differences are that we are primarily concerned with private goods, and we rule out the use of an external medium of exchange whereas the literature just mentioned relies heavily on such a medium. Two approaches are of interest: (a) The "demand-revealing procedure," presented in Tideman and Tullock (1976) and in Tideman (1977). The former provides a lucid review of earlier works and develops a generalized procedure oriented toward practical implementation; the latter contains detailed discussion of several aspects of the procedure. The demand-revealing procedure permits dominant strategies when preferences are linear in the external medium of exchange; hence each individual has a "canonical response" which depends only on his own preferences. Strictly speaking, the procedure is not efficient; it includes an "incentive tax" (over and above taxes necessary to finance the public goods), the proceeds of which must be wasted in order not to distort the incentives. (b) The "optimal government" of Groves and Ledyard (1977). This procedure is efficient for quite general preferences, but one individual's response will in general depend on everybody's preferences; the solution is a Nash equilibrium and dominant strategies need not exist.

⁶ Our interest in this problem arose partly in response to a case in which there was a practical need for such a procedure: the assignment of freshmen to upperclass houses (residence halls) at Harvard College. In 1975 and 1976, a rank-ordering procedure was employed (see n. 13 below and accompanying text). It was observed that a relatively small number of students were assigned to their first choice. Since it was thought that the students had a much stronger preference for first over second place than for second over third, a procedure was introduced in 1977 that gives priority to first preferences. In this procedure the individuals are ordered by lot and are then assigned in order to their first-place houses as long as slots are available. An individual whose first-place choice is filled is skipped. Only when first-place choices are fully allocated does the system come back to second place, then back to third place, etc. Unfortunately, this system generates strong incentives for strategic behavior. For example, a student might not list his true first-place house if he thought it would be listed first by many others. By listing his less popular second choice first he would give up a small chance of getting his first choice for a significant increase in his chance of getting his second choice. Harvard administrators believe that many students did in fact behave strategically in 1977.

Though some of the procedures required to demonstrate the success of the allocation mechanism can best be described with the aid of mathematics, the essential logic supporting our results can be understood intuitively. The more technical aspects of our analysis are presented in four appendices.

I. Formulation of the Problem

Formally, the structure of the problem is as follows: Each of I individuals must be assigned to one and only one of J jobs. Giving individuals lotteries over jobs is permissible, as long as the lotteries are ultimately resolved so as to yield definite assignments. It is required that exactly M_j individuals be assigned to job j , for $j = 1, \dots, J$. The numbers M_1, \dots, M_J are given and sum to I .

Any individual i has a given utility u_{ij} for being assigned to any job j . If i is assigned probabilistically to one of several jobs, i 's ex ante evaluation of this probabilistic assignment is i 's expected utility for that lottery on jobs. In other words, u_{ij} is i 's von Neumann-Morgenstern utility assessment for job j .⁷ It is not necessary to determine any origin or unit of measurement for these utility scales; all results are unique for any positive linear transformation of them. We do not assume that i would automatically truthfully disclose his u_{ij} value for job j . Indeed, we assume that he engages in strategic behavior to maximize his expected utility, given his expectations about the utilities that other individuals will express and the way the procedure will process those expressions in conjunction with his own to assign him a job or lottery on jobs.

We have made three simplifying assumptions regarding the structure of preferences. First, each person's preferences are assumed to concern solely his own assignment; he is indifferent to the assignments of others.⁸ In addition, jobs are assumed not to have preferences for the persons assigned to them; that is, this is not the well-known marriage problem.⁹ Finally, we assume that values external

⁷ Hence we assume that the individuals satisfy the behavioristic assumptions that are necessary to prove that preferences over lotteries can be represented by a utility function. See, for example, Luce and Raiffa (1957, chap. 2), where a sufficient set of such conditions is presented and the realism of the conditions discussed.

⁸ In a limited way it is possible to allow preferences to apply to the assignment of others. Consider freshmen who have preferences not only among upperclass houses but also for individuals who will live there with them. If their number is small relative to the capacity of the house, it will be possible to treat a group of students as a unit so long as the exact capacities can be met. Usually this will be possible if most individuals are applying singly, or if there is sufficient flexibility to allow houses to be filled a few students above or below quota.

⁹ The case of bilateral preferences (the marriage problem) is presented in Gale and Shapley (1962). The problem of honest revelation of preferences is not considered in that paper.

to the individuals' preferences play no role. Hence an institutional interest on the part of a national organization to have all regions of the country represented on all delegates' committees, will not be taken into account.¹⁰ Clearly, these three restrictions may prove to be important in some practical circumstances. (The extent to which they can be relaxed is an area for further study.)

For the moment we are requiring that each person be assigned to exactly one position and that the capacities of the positions, given by the numbers M_j , be exactly filled. These capacity restrictions merely facilitate exposition and can be relaxed (see Appendix D).

The problem can be written formally as follows. *Given:* I, J, M_1, \dots, M_j , where $\sum_j M_j = I$, $\mathbf{U} = (u_{ij}), i = 1, \dots, I; j = 1, \dots, J$; that is, $\mathbf{U} = (\mathbf{u}_1, \dots, \mathbf{u}_i, \dots, \mathbf{u}_I)$, $\mathbf{u}_i = (u_{i1}, \dots, u_{iJ})$.

Each i submits a preference statement, $\mathbf{w}_i = (w_{i1}, \dots, w_{ij}, \dots, w_{iJ})$, which may be strategic rather than truthful. Thus the assignment procedure receives a total submission $\mathbf{W} = (\mathbf{w}_1, \dots, \mathbf{w}_i, \dots, \mathbf{w}_I)$. It then generates allocations meeting both the constraints on the M_j 's and the requirement that each individual be placed in precisely one job.

The objective is to find an assignment procedure that in a general range of circumstances will elicit a \mathbf{W} and process it to generate an efficient allocation, one that is Pareto optimal. Evaluations of efficiency, it should be stressed, are made in terms of \mathbf{U} , not \mathbf{W} . (Only in the case where all individuals provide truthful preferences will the two be identical.)

The procedure we shall discuss allows individuals' preferences to be weighted in a variety of ways. Unless the weightings are to be very unequal, or preference orderings are remarkably dissimilar, no individual can be assured his first-choice job.¹¹ This implies that, at least at some stage in the process, individuals will be assigned to lotteries over jobs. That is, on the basis of \mathbf{W} , individual i will be assigned to jobs with the probabilities $\mathbf{p}_i = (p_{i1}, \dots, p_{ij}, \dots, p_{iJ})$. The total assignment

¹⁰ Such an institutional interest could be included as a system of constraints. In particular, if the rules specify exactly the number of members from each region on every committee, one can simply perform a separate assignment for each region. If the rules place maxima and/or minima on representation on the committees from the regions, an appropriate generalization of our procedure can presumably be constructed. But if there are many dimensions of such constraints, the procedure becomes unwieldy. And in any case the institutional interest must be represented as constraints; it cannot simply be traded off against the individuals' preferences.

¹¹ If the individuals are ordered in priority, then Pareto optimality can be assured by allowing individuals to select in turn their preferred job among those still available. This is an extreme case of unequal weighting and can only be achieved as a limiting case of our procedure. We believe that in general lexicographic ordering of the individuals' utilities will not be regarded as equitable, and we therefore do not consider the case. Certain organizations may, however, have systems that lead to such structures, as, for example, when employees bid on positions according to strict seniority.

is thus $\mathbf{P} = (\mathbf{p}_1, \dots, \mathbf{p}_I)$. The constraint that each individual with certainty be assigned to precisely one position requires that

$$\sum_j p_{ij} = 1 \quad \text{for all } i. \quad (1)$$

In order that the capacities of the jobs be filled exactly, it is necessary that the expected number of individuals assigned to a job equal its capacity, that is, that

$$\sum_i p_{ij} = M_j \quad \text{for all } j. \quad (2)$$

Ultimately, the lotteries represented by \mathbf{P} will have to be simultaneously resolved in a manner that assigns each individual to one job, precisely meets the capacity constraints, and gives everybody the right probabilities of being assigned to the various jobs. It turns out that (1) and (2), together with the nonnegativity of the numbers p_{ij} , is sufficient to guarantee that such a final assignment procedure can be implemented. Straightforward mechanisms for resolving the lotteries will fail. For a successful procedure see Section III and Appendix C below.

Ex Post and Ex Ante Pareto Optimality

Given that the procedure assigns individuals to lotteries over jobs, there are two possible interpretations of the Pareto optimality condition. The procedure is Pareto optimal *ex ante* if its assignments of lotteries to individuals is Pareto optimal relative to their preferences among lotteries. The procedure is Pareto optimal *ex post* if the final assignment of positions to individuals produced by performing the lotteries is Pareto optimal relative to their preferences among positions. It is readily shown that if a procedure is Pareto optimal *ex ante*, then it must be Pareto optimal *ex post*.¹²

The converse need not be true. Consider three individuals with utilities for jobs A, B, and C that are, respectively, 100, 10, 0; 100, 10, 0; and 100, 80, 0. An assignment procedure that selects an individual

¹² Proof: Let a lottery assignment be given, that is, assume nonnegative numbers p_{ij} satisfying (1) and (2) are given. Also assume that this lottery gives a final assignment which is not efficient. Then there exist two sets of job assignments j_1, \dots, j_I and j'_1, \dots, j'_I , satisfying the capacity constraints, such that (i) in the given, final assignment, individual 1 gets job j_1 , 2 gets j_2 , etc.; (ii) If instead 1 were assigned to j'_1 , 2 to j'_2 , etc., nobody would be worse off and somebody would be better off. Let a be the smallest of the numbers $p_{1j_1}, \dots, p_{Ij_I}$. By (i), a is positive. Then reduce each of the numbers $p_{1j_1}, \dots, p_{Ij_I}$ by a and add a to each of $p_{1j'_1}, \dots, p_{Ij'_I}$. Since both assignments satisfy the capacity constraints, (1) and (2) still hold, and the choice of a guarantees $p_{ij} \geq 0$. Therefore, the new assignment of probabilities represents a feasible lottery (see Appendix C). From (ii) and the nature of the preferences we can conclude that this change has made somebody better off without making anybody worse off. Hence the lottery assignment was not *ex ante* optimal, and the proof is complete.

at random and offers him his first choice, then selects one of the remaining two at random and offers him his choice, will result in each individual getting the same lottery on positions ($\frac{1}{3}$, $\frac{1}{3}$, $\frac{1}{3}$). Honest preferences will be provided and assignments will be Pareto optimal ex post (as indeed all assignments must be when ordinal rankings for positions are identical).¹³ However, the assignment procedure is not Pareto optimal ex ante. The individuals' expected utilities are $36\frac{2}{3}$, $36\frac{2}{3}$, and 60, respectively. All would do better if the third individual were assigned to job B and the first two tossed a coin to see who got job A. Expected utilities would then rise to 50, 50, and 80. The initial set of lotteries is thus shown to be dominated.

Our efficiency requirement then is to achieve ex ante Pareto optimality, which in turn guarantees ex post Pareto optimality.

II. A Procedure to Achieve an Efficient Allocation: A Pseudomarket for Probability Shares

Market mechanisms are the traditional means to produce efficient allocations of private goods. It is not surprising, therefore, that our allocation procedure employs a market-like process. In theory, the entire allocation could be conducted with individuals constructing their lotteries by purchasing probability shares in the positions on a decentralized basis, with a central market maker who announces prices and then adjusts them in response to excess demands. In practice, it turns out that the procedure works much more simply and effectively by asking individuals to provide their preferences and then conducting the market on a simulated basis. This is not possible with ordinary markets, for a full description of an individual's preferences would be too unwieldy. When the purchased commodity is a lottery, however, and the preferences for outcomes are expressed by von Neumann–Morgenstern utilities, all that is required is that each individual provide a vector of J numbers that gives his utility value for each of the J positions. Thus in a three-job world, individual i 's preferences might be represented 0, 100, 42. Asking for preferences rather than selling probability shares directly will save individu-

¹³ For any set of individuals' preferences, a procedure similar to the one described here will achieve ex post Pareto optimality and elicit honest revelation of preferences. When there are no ties in anybody's preference ordering on positions, one can proceed exactly as described in the text: order the individuals by lottery and let them choose. When the possibility of ties cannot be ruled out, a somewhat more complicated procedure is necessary to guarantee Pareto optimality. Still it will be impossible for any one person to gain by misrepresentation of preferences. But it may be possible for a group of persons, by coordinated misrepresentation, to achieve a change from which some of them win and the rest do not lose. No deterministic procedure can rule out this possibility and at the same time guarantee Pareto optimality.

als considerable time and effort since they need not find their optimal lotteries over each of a large number of iterations of prices.

Still, it is necessary to have people list their utilities. In practical applications, one cannot expect the participating individuals to be familiar with the concept of a utility function. Care must therefore be taken to formulate questions in such a way that people understand and respond correctly. We do not believe this is an insurmountable task. There is a danger, however, that people will not find it worthwhile to spend a lot of time to think out and report accurate and detailed preferences. (They may, for example, lump together everything but the three positions they rank highest.) Compared with accurate reporting, this may lead to some loss in allocational efficiency, but the loss is not likely to be significant, given individuals' relative indifference.

For clarity, we first illustrate the working of the procedure and then turn to examine other critical issues, including most significantly the way market-clearing prices can be determined, whether individuals will provide honest preferences, and how initial endowments in the pseudomarket should be determined.

III. The Procedure at Work

Let us assume that everyone has submitted his honest von Neumann–Morgenstern utility vector to the assignment procedure and that budgets have been established. (See Section IV below.) The next step is to efficiently assign individuals lotteries, given the constraints on capacities for each job. This is done by simulating a market. The “commodities” that are sold are probability shares in positions. These commodities are infinitely divisible. An individual purchases—or rather, the central mechanism purchases on his behalf—his preferred lottery given the prices for probability shares in positions, his budget constraint, and the requirement that the sum of his purchased probabilities must equal 1, that is, that he will certainly be assigned to one position. (To avoid circumlocutions, we shall refer below to an individual purchasing for himself when in reality the mechanism would be his agent and do the purchasing for him.)

To describe the process more precisely, let B_i be a positive number representing the budget for person i , and assume that the vector $\mathbf{q} = (q_1, \dots, q_j, \dots, q_j)$ is given, where q_j , a nonnegative number, is the price of probability shares in position j . (A 0.4 chance at position j would thus cost $0.4q_j$.) Then person i shall choose numbers p_{i1}, \dots, p_{iJ} to maximize

$$\left. \begin{array}{l} EU_i = \sum_j p_{ij} u_{ij}, \\ \text{subject to} \\ 0 \leq p_{ij} \leq 1 \text{ for all } j, \\ \sum_j p_{ij} q_j \leq B_i, \end{array} \right\} \quad (3)$$

and constraint (1). Since the individual's objective function is linear in his purchased probability shares, he would buy into only a single position were it not for the constraint that the numbers p_{ij} sum to 1. In cases where both this constraint and his budget constraint are binding, there exists a solution to the optimization problem in which shares in precisely two positions are purchased.

The utility maximization problem for an individual can have many solutions. In order to guarantee Pareto optimality, we require that whenever an individual has a choice among several lotteries with the same expected utility, the least expensive of these shall be chosen. (Still the solution need not be unique.) This rule does not harm the person who has a nonunique solution to (3); by definition, two lotteries are equally good when their utility is equal. But the rule can benefit others. (It is easy to understand intuitively why this rule is reasonable: When you do not care which of two lotteries you get, you help others by staying away from the one that is more highly valued in the market.)

When all individuals have made their choice, we want the total demand for probability shares in the positions to be exactly $M_1, \dots, M_j, \dots, M_j$, that is, we want (2) to be satisfied. In general this will not happen; there will be excess demand or supply for some of the positions. The assignment mechanism then responds by adjusting prices. The critical first question is: Do there exist market-clearing prices? The answer is yes. There will always exist a price vector \mathbf{q} such that there exist numbers \mathbf{P} chosen to maximize individuals' expected utilities in accordance with the description above and satisfying the constraint (2) on capacities.¹⁴ The market-clearing price vector can

¹⁴ Note the double use of existence; it is not necessarily the case that the price vector \mathbf{q} leads unequivocally to market clearing. One can construct examples with the following properties: There is a unique price vector \mathbf{q} that can clear the market. For this price vector, individual 1 can choose among an infinity of probability vectors with equal maximal utility and equal minimal price. But in order to clear the market, person 1 must choose a particular vector from this infinite set. To decide which vector person 1 shall choose, we need information about the preferences of persons 2, 3, etc.

In the market-clearing assignment, there may be persons who are assigned positive probability shares in more than two positions, although, as noted above, every person's individual optimization problem has a solution in which at most two positions are bought.

The existence of market-clearing prices is proved in Appendix A, from purely

always be chosen so that at least one price is 0. Hence we need only consider such price vectors in the first place. This has the advantage that there is never any doubt that each individual can buy a lottery which satisfies (1); whenever the budget runs short, one can fill up with probability shares in a free position.

Pareto Optimality of the Cleared Market

Having the market clear in the manner we describe guarantees that the assignment procedure leads to an outcome that is ex ante Pareto optimal. To prove this, let \mathbf{p}_i be the lottery over positions that individual i receives in the assignment \mathbf{P} , which results from the market-clearing price vector \mathbf{q} . Assume that \mathbf{P}' represents another assignment which is better for somebody and not worse for anybody. Let i_1 be a person who is better off in the latter assignment, with his lottery over assignments being \mathbf{p}'_{i_1} . Then we must have $\mathbf{p}'_{i_1} \cdot \mathbf{q} > B_{i_1}$, and $\mathbf{p}'_{i_1} \cdot \mathbf{q} > \mathbf{p}_{i_1} \cdot \mathbf{q}$; otherwise person i_1 did not choose an optimal lottery originally. Since $\sum_i \mathbf{p}_i = (M_1, \dots, M_J) = \sum_i \mathbf{p}'_i$, there must exist a person i_2 such that $\mathbf{p}'_{i_2} \cdot \mathbf{q} < \mathbf{p}_{i_2} \cdot \mathbf{q}$. Then i_2 must be worse off with the lottery \mathbf{p}'_{i_2} than with \mathbf{p}_{i_2} ; otherwise i_2 would not have chosen \mathbf{p}_{i_2} in the first place. (Here it is essential that the least-expensive lottery be chosen if utility is equal.)

By an earlier result, the described procedure also leads to an assignment that meets the less restrictive condition of being ex post Pareto optimal.

mathematical principles. Presumably, the result can also be obtained as a special case of general equilibrium results. The subsumption is not immediate. The standard results either place restrictions on the consumption set which are not met in our model because of (1), or they impose assumptions which are not satisfied in our case in order to rule out certain "exceptional cases." (See the review of results in Quirk and Saposnik [1968].) There is, moreover, a difference between our model and one used frequently in general equilibrium results. We specify the initial endowments in the form of real numbers, B_i . (The scale of these numbers is, of course, arbitrary. Multiplying each B_i by a positive number will just lead to the equilibrium prices being multiplied by the same constant.) Often, individuals are assumed to start out with initial bundles of goods. If the latter formulation is used in our problem, equilibrium prices need not exist. Example: There are three individuals and two positions; position A with capacity 1 and position B with capacity 2. Individuals 1 and 2 prefer A to B; 3 prefers B to A. (With only two positions, ordinal preferences on positions uniquely determine utility schemes.) We want to treat the three persons equally. In the framework of this paper, we set $B_1 = B_2 = B_3 = 1$. Market-clearing prices are 2 for A, 0 for B, and corresponding lottery assignments are $(\frac{1}{2}, \frac{1}{2})$ for 1 and 2, (0, 1) for 3. In the alternative framework, one can treat everybody equally by giving each an initial endowment $(\frac{1}{3}, \frac{2}{3})$. But then there are no equilibrium prices. (If one insists on expressing the initial endowment in terms of bundles of commodities, this problem can presumably be solved by introducing some redistribution of endowments in cases where a person wants to buy a bundle which is less expensive than the initial endowment, as is the case for C in the example.)

TABLE 1
UTILITIES FOR POSITIONS

INDIVIDUAL	POSITION			INDIVIDUAL	POSITION		
	A	B	C		A	B	C
1	100	86	0	12	80	100	0
2	100	71	0	13	60	100	0
3	100	57	0	14	40	100	0
4	100	43	0	15	20	100	0
5	100	29	0	16	0	100	75
6	100	14	0	17	0	100	50
7	100	0	83	18	0	100	25
8	100	0	67	19	67	0	100
9	100	0	50	20	33	0	100
10	100	0	33	21	0	50	100
11	100	0	17				

An Example

Consider a world with 21 individuals and three positions with seven places each; call them A, B, and C. The group as a whole has a tendency to prefer A to B to C. There are six people with ordering A, B, C; five with A, C, B; four with B, A, C; three with B, C, A; two with C, A, B; and finally one with C, B, A. The von Neumann–Morgenstern utilities of the individuals for the positions are shown in table 1. Giving each individual a budget of 1, the market-clearing prices in this example are $q_A = \frac{9}{7}$, $q_B = \frac{8}{7}$, and $q_C = 0$. At these prices, individuals 1 and 12–18 will demand the probability vector $(0, \frac{7}{8}, \frac{1}{8})$; individuals 2–11 will demand the probability vector $(\frac{7}{10}, 0, \frac{3}{10})$; and individuals 19–21 will demand the probability vector $(0, 0, 1)$. Everyone has a unique optimal lottery assignment. The individuals' budget constraints are met; the available positions are just exhausted.

Market-clearing Prices

In the simple example above, prices to clear the market were readily computed. How can we compute market-clearing prices in general? This turns out not to be a trivial question. The simple iterative process of starting with one price vector and adjusting prices up or down according to whether there is excess demand or excess supply will not necessarily converge. One problem is that constraint (1) is likely to create Giffen goods; then an increase in the price of a commodity for which there is excess demand can actually increase the demand. In a variety of examples that we have considered, we have always been able

to compute market-clearing prices through a simple trial-and-error procedure. We believe that trial-and-error methods along this line will always work.

Still, we should like to have an algorithm that can be demonstrated to give a satisfactory result in all cases. Thus we turn to the pioneering work of Herbert Scarf. Our adaptation of his algorithm will, within a workable time period, enable us to compute a price vector and an assignment of probability shares to persons, such that the market is cleared and everybody's assignment is arbitrarily close to being optimal at the given prices. Moreover, the allocation will yield an outcome that is arbitrarily close to the Pareto frontier. (The prices need not be close to a true market-clearing price vector, however.) For all practical purposes this should be fine. The algorithm and the way it deals with certain complexities are discussed in Appendix B. Where necessary below, we assume that we can compute market-clearing prices and lottery assignments.¹⁵

Conduct of the Lottery

Hypothetical budgets have been passed out, preferences have been stated, and the pseudomarket has been run and has assigned lotteries to individuals. The final step in the assignment procedure is to conduct the lottery so as to just fill the spaces in a manner that offers each individual precisely the lottery on positions that he has purchased. This turns out to be surprisingly difficult. Obviously, it is not possible to run each individual's lottery separately, since chance could overflow and underfill positions. Running individuals' lotteries in a sequence is a logical alternative; it is a procedure that fails less obviously. After one individual's lottery is run, the probabilities in the lotteries for "unresolved" individuals must be adjusted in order to meet the capacity constraints. However, the adjustment for the second individual depends not only on the outcome of the lottery for individual 1 but also on the way the lotteries for all subsequent individuals divide up the remaining capacity. The adjustment procedure becomes unwieldy, for it must look ahead all the way to the end.¹⁶

¹⁵ We mentioned in the introduction that if utilities are known and weights assigned to the individuals, our problem would be equivalent to the job assignment problem. When a solution \mathbf{P} is found by our method, it is always possible afterward to assign weights v_1, \dots, v_I to the individuals such that \mathbf{P} maximizes $\sum_i v_i \sum_j p_{ij} u_{ij}$ over all feasible \mathbf{P} , and every final assignment of individuals to positions that can result from \mathbf{P} is a solution to the corresponding job assignment problem. But the weights depend on people's utilities. And if we somehow fixed the weights and announced that the job-assignment algorithm would be used, there would be strong incentives to misrepresent preferences.

¹⁶ To prove that it is necessary to look ahead, consider an example based on fig. 1. If person 1 gets position A, what must happen to person 2's probability of getting position

Fortunately, a simpler algorithm is available. Consider the matrix \mathbf{P} of probability shares ("purchased" on the pseudomarket). A representative \mathbf{P} matrix is shown in figure 1.¹⁷ First eliminate from the matrix any row which contains only one positive number. (This positive number must then be 1, and it therefore determines the position to which the corresponding person shall be assigned.) In the reduced matrix, find a cycle of positive numbers. Start at any positive entry and move horizontally until another positive entry is found. (This is always possible since all rows in the reduced matrix have more than one positive entry.) Then move vertically to another positive entry. (Again, this always exists. The sum in each column is an integer, so if there were only one positive entry in the column it would have to be 1. But then the corresponding row should have been eliminated in the previous step.) Continue to alternate between horizontal and vertical movements until the path comes back to a row or a column in which it has been before; then close the cycle in the next step by going to the previously visited entry in that row or column. (The cycle thus constructed need not contain the entire path we have traced out; there may be an initial segment which does not belong to the cycle and should be discarded.) This will always work, because the matrix is finite; for a formal proof, see Appendix C.

In the example below, there are many cycles. The simplest ones consist of only four points. To illustrate the subsequent discussion, a slightly more complicated one is indicated by circles and squares.

Once the cycle is found, its points are divided into two classes; moving along the path, alternate points are placed in the two different classes. In the figure, one class is indicated by squares and the other by circles. Now we shall choose one of the classes by lottery and reduce all members of that class and increase all members of the other class, all increments being of the same absolute value. This will not

C? If 1 and 2 get A and C, then person 3 must get B or D. The table indicates that the initial probability for this B or D outcome is .4. Thus the parley 1 in A, 2 in C can have probability no greater than .4. Since 1 has the probability of .7 of getting A, then the conditional probability of 2 getting C when 1 gets A cannot exceed $\frac{2}{7}$. Assume, however, that individuals 3 and 4 had different initial probability assignments so that the last two lines in the matrix were

3	.3	0	0	.7
4	0	.7	.3	0

If 1 now gets A, the probability of 2 getting D must be made equal to 0, otherwise 3 would get a positive probability of getting B or C, contrary to assumptions. But 2's probabilities of getting C and D must sum to 1. Hence the adjustment of 2's probabilities when 1 gets A, must depend on the data for persons 3 and 4.

¹⁷ It will generally be the case that some persons are assigned positive shares in more than two positions, as is the case for person 3 in the example (see second paragraph of n. 14 above).

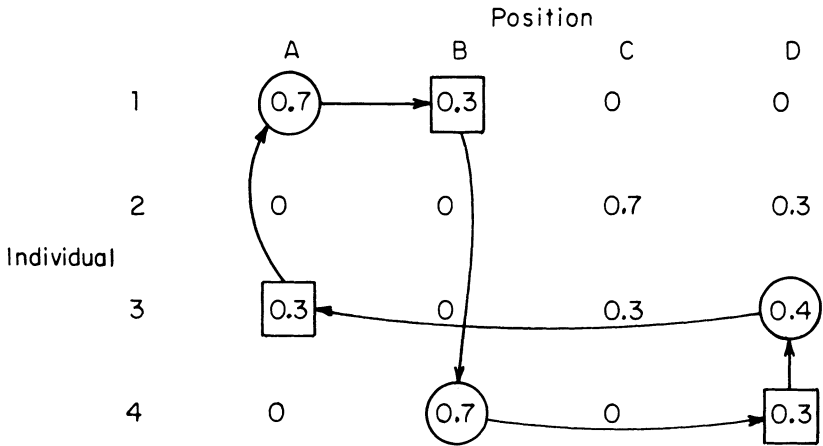


FIG. 1.—A probability cycle in the assignment procedure

change the sum in any row or column, since each row or column contains one point from each class or no points at all from the cycle. The increment shall be such that at least one positive element is reduced to zero and no element is reduced below that. In the example, if the “circle points” are to be reduced, the increment shall be 0.40; if the “square points” are to be reduced, the increment shall be 0.30.

Moreover, the expected increment of every element in the matrix shall be zero; the probability of each of the classes being chosen for reduction is computed so as to achieve this. In the illustration, the circle points shall be reduced with probability $\frac{3}{7}$; the square points with probability $\frac{4}{7}$. Since $\frac{3}{7} \times 0.4 = \frac{4}{7} \times 0.3$, this gives an expected increment of 0.

When this step is carried out, the number of positive numbers in the matrix is reduced. If necessary, the step is performed again and again and again, but it must stop in a finite number of steps. Individuals are assigned to positions with probabilities that are consistent with their initial lottery assignments. A more detailed and formal description of the procedure is given in Appendix C.¹⁸

¹⁸ Comments received on a preliminary draft of this paper have made us aware that the algorithm described above is closely related to earlier works. For one thing, the procedure is similar to the well-known “stepping-stone” algorithm for the transportation problem in linear programming. More directly related to our problem, the results that (1), (2), and $p_{ij} \geq 0$ are sufficient to guarantee that the lottery can be carried out and the capacities filled was proved in Birkhoff (1946). A proof of the same, which is similar to our algorithm, is given in von Neumann (1953). (Birkhoff and von Neumann consider the slightly less general case of $I = J$, but the generalization is immediate.) We thank Elon Kohlberg for bringing the literature on this result to our attention.

Honest Provision of Preferences

We have demonstrated that if we know individuals' von Neumann–Morgenstern utilities we can generate an assignment matrix that is ex ante Pareto optimal and that we can then resolve that matrix to produce an assignment of individuals to jobs. Can we be confident, however, that individuals will provide honest preferences? The answer is yes, as long as two conditions are met. First, the assignment mechanism must employ these preferences to choose for the individual as he would choose for himself. That is, it must select his optimal lottery subject to market prices, his budget constraint, and the constraint that his probability shares must sum to one. Second, no individual i should have a noticeable ability to manipulate market prices in a manner favorable to himself through the manipulation of his expressed preferences, that is by submitting $\mathbf{w}_i \neq \mathbf{u}_i$.

Our proposed procedure meets the first condition; it acts honestly as an agent. What about the second, the possibility of distortion affecting prices in a favorable manner? Distortion may have two effects on the welfare of the distorting individual. First, he may lose by receiving an inefficient bundle for whatever prices are established. Second, he may gain by securing a more favorable vector of prices. Distortion can be expected when the gain to a potential distorter outweighs his loss. For whom might this be the case?

Consider an individual who is *influential* in the sense that small changes in his expressed preferences can alter prices. This can only be the case for a person who is perceived by the system as being indifferent among a continuum of lotteries, so that his ultimate assignment will be chosen in order to secure a cleared market.¹⁹ If he misrepresents his preferences prices will change, a number of noninfluential people will change their purchase, and the allotment of the influential person will be adjusted. This change may or may not be to his advantage, depending on the total number of persons, their preferences, and the nature of the misrepresentation.

Let us first examine a situation in which misrepresentation should be expected. There are three positions, each with capacity 1, and three individuals, with identical endowments and the utilities 100, 90, 0; 100, 60, 0; and 100, 10, 0. The second individual is influential. If all individuals express their true preferences, his expected utility would be 53.3. By overstating his utility for the second position, the second person can, however, increase his utility. His optimal representation is 100, 77.46, 0, which gives him a utility of 55.0.

¹⁹ See first paragraph of n. 14 above. If we rule out cases that essentially have probability 0, such as two persons having exactly the same utilities in a world with three or more jobs, the number of influential persons is at most $J - 2$. It may be lower, even 0.

Such misrepresentation will necessarily hurt an individual if it moves him out of the area where he is influential. If there are many persons with diverse preferences the range over which one individual is influential will be quite small. The gains from any misrepresentation even in this region will be minimal and probably negative. An individual who is not influential can only gain by moving into a region in which he will be influential, and, if numbers are large, chances of gain are slight. Unless an individual has information revealing the range over which he will be influential, information unlikely to be available in any real world circumstance, there will be at best a highly unfavorable gamble involved in distorting one's preferences.

Consider a large numbers case where any misrepresentation would lead to losses. To extend our previous example, say that each position has capacity 27 and there are 81 individuals with utility vectors 100, 90, 0; 100, 89, 0; 100, 88, 0; . . . ; 100, 10, 0. The individual whose preferences are 100, 60, 0 is influential. Now, however, he would lose were he to misrepresent his utilities in any way.

Our experience with the three-good case leads us to conjecture that when no individual can gain through misrepresentation of preferences no coalition of individuals can gain either.²⁰ Note that a coalition attempting such misrepresentation would be handicapped by its inability to make side payments; Pareto improvements would have to be sought. If our conjecture proves correct, at least with respect to its protection against strategic action by a coalition, our pseudomarket has a security advantage over traditional markets, in which side payments are possible.

In our model, as in traditional markets, no individual can have a foreseeable effect on price if there is a sufficient number of participants. (Indeed, even with omniscience no individual may be able to distort preferences to his benefit.) Honest preferences will be provided. In this instance, the number of individuals must be measured relative to the number of positions; many individuals have to be interested in each. In traditional markets, similarly, nonstrategic behavior requires that there be many buyers and sellers for each commodity.²¹

²⁰ Here we assume that no two individuals happen to have the same utilities. If there is a group of people with exactly the same preferences, it is possible that the group can be "collectively influential" and gain, in an expected-value sense, by coordinated misrepresentation, although no single member who unilaterally misrepresents preferences can gain.

²¹ Nothing in our discussion suggests that the market-clearing price vector is unique; indeed, it is easy to construct examples where uniqueness does not hold. It may matter for the individuals which market-clearing price vector is chosen; some may be better off in one situation than in the other. (Since the allocations are Pareto optimal, the

IV. Equity, Endowments, and the Absence of a Medium of Exchange

A prime motivation for considering this problem was its empirical importance. The efficient allocation of individuals to positions is sought in a wide range of situations in which no external medium of exchange is employed. The recommendation that introducing money would handle all problems is unlikely to be helpful. The use of any currency external to the particular allocation problem is generally unacceptable because the loss in equity, or perceived equity, is thought to more than outweigh any gain in efficiency. (For example, few university faculties would allow their wealthier members to buy their way out of committee assignments; see n. 2 above.)

Still, accepting the prohibition on side payments, there will be a large number of efficient allocations. How should we choose among them? If equal treatment is a goal, as it would be in many situations, then it would seem that, in our procedure, individuals should start with identical initial endowments or budgets.²² Equal budgets will not assure equally good outcomes for all, nor is such a requirement a reasonable one. If one person, for instance, prefers a position that generally is unpopular, it hardly seems inequitable to guarantee this person that position, despite the fact that no one else gets more than a 60 percent chance at his first choice.

The precise consequence of giving everybody the same budget is to give each the same opportunity set. That is, for any price vector the set of lotteries from which an individual shall choose an optimal one—the set defined by (1) and (3)—is the same for all persons. This, in turn, implies that in the final assignment no person will prefer somebody else's lottery assignment to his own. (No one will envy his neighbor before the lottery is conducted.)

In many cases equal treatment may not be desired. We may want to favor some and discriminate against others while giving equal treatment within certain groups, all of this being done independently of

differences must go both ways.) This suggests another type of strategic behavior, namely, to misrepresent one's preferences in such a way that the system will end up at a different market-clearing price vector than it would otherwise. But this will require even more information than the other types of strategic behavior; in particular, one must be able to simulate the working of the algorithm which computes the price vector. It is also possible to build random elements into the algorithm, so that it will be impossible to predict which market-clearing price vector it will reach when there is more than one of them.

²² An alternative, which perhaps would make the equity of the situation more obvious, is to give each individual an initial endowment of M/I probability share in each position j and then let him trade these initial endowments. As noted at the end of n. 14 above, the existence of equilibrium prices is not guaranteed in this case, at least not unless some complicated redistributive procedure is employed.

individuals' preferences. Thus seniority confers preference in committee assignments in most legislatures, upperclassmen may be assigned priority in the choice of dormitories, etc. Urban homesteading procedures may be designed to favor members of groups that have been traditionally disadvantaged in housing markets or, in a quite different spirit, to help those who have attributes that suggest that they will make good use of their opportunity. Still, it is not easy to say what favored should mean, but it certainly implies that if i_1 is to be favored compared with i_2 , then the former shall get an assignment which, in i_1 's own eyes, is at least as good as the one i_2 gets. A larger budget endowment for i_1 will achieve this.²³

In most practical applications, we believe, equal budget endowments would be employed, for they provide an exceedingly powerful signal of both equity and equality. How to establish budget disparities when equality is not desired is a matter meriting future exploration.

V. Extensions and Generalizations

The model that we have considered can be extended in many ways. Some such extensions would be required to make the method applicable to a wide range of policy issues. For example, a number of assignment lotteries would have to be conducted simultaneously if it were desirable to assign groups of individuals to the same position. This complication might arise if the individuals were students wishing to go to the same dormitory (see n. 8), or if the "jobs" were quantities of energy being allocated to different plants where it would be infeasible to have 30 percent of a plant's capacity met by coal and the rest by natural gas.²⁴

²³ Another question is, How much better off is i_1 than i_2 if he has, say, twice as high a budget? We merely state this question and make no attempt to answer it. A particular form of unequal treatment is to give one group of individuals absolute priority over another. (If each "group" here is but one person, we are back in the lexicographic ordering; see n. 11 above.) An example can be course assignment in colleges; we may want to give seniors an absolute priority. In our procedure, this can be achieved as a limiting case, with infinitely different budgets. More easily, one can use a two-stage procedure. First, only the high-priority group takes part, and the capacities on the positions are treated as upper bounds rather than exact capacities. (This is possible; see Appendix D.) Second, the assignment is made for the low-priority group, and the capacities are whatever is left over.

²⁴ Our original interest in the simultaneous-assignment problem arose in connection with U.S. energy policy. There may be limited supplies of various fuels in different geographic areas or in the nation as a whole. Given the general refusal to let the market allocate fuels, it is now being proposed that government permission be required to install new types of facilities (for example, gas boilers for heating), or even to keep old ones. The potential inefficiencies inherent in such a centralized system of allocating fuels are enormous. An approach in the spirit of this paper might at least avoid some of the problems. The need for joint assignments, hence simultaneous resolution, arises because different plants will have different requirement levels. Plant A may use 100 million btus per period, whereas plant B requires 300 million. If a single boiler is to supply all btus, obviously all the btus for a plant will have to be assigned together.

Appendix D is directed to generalizations of our methods. It first shows that it is not necessary to have the total capacity of the positions exactly equal to the number of persons. In particular, it discusses procedures that are appropriate when there is just an upper and/or lower bound on the capacities of positions.

An important generalization enables us to give an individual a number of assignments. Several classes of cases are discussed; they correspond to various conditions imposed on individuals' preferences when multiple assignments are made. An intriguing application relates to fair division problems. A study of indivisible items is given. They must be divided equally among a group of people whose preferences may vary. Our method allows limited complementarity of preferences, up to the situation in which an individual's tradeoff rate between additional units of A and B can depend on the quantities of A and B he already possesses, though not on quantities of C and D. Moreover, it permits alternate criteria for equality. An obvious criterion would be that each individual start with the same endowment, but other possibilities exist (see Section IV above).

Interesting further work, we believe, will not only extend our methodology but will explore the challenging problems that will arise as these methods are applied in real-world contexts, a development we hope to foster.

VI. Conclusion

The efficient allocation of individuals to scarce positions involves a four-step procedure. First, they are given hypothetical endowments reflecting their relative strengths of claim for positions. (A significant special case offers equal endowments.) Second, individuals' von Neumann–Morgenstern utilities for the alternative positions are elicited. Third, a pseudomarket is employed to assign lottery shares to each individual to produce an efficient outcome. Fourth, a specialized mechanism is employed to conduct the lottery. It is noteworthy that the procedure leads to an outcome that is Pareto optimal both *ex ante*, before the lottery is conducted, and *ex post*, assuming that the expressed utility values are honestly provided. Since the procedure operates what is in effect a pseudomarket, individuals will in general reveal their utilities honestly as long as the number of individuals is large relative to the number of positions.

Increasingly, our society is choosing to allocate goods and services on the basis of perceived need or entitlement rather than through the market. Generally, the underlying philosophical argument is that for the good under consideration wealth, or other considerations external to the problem at hand, should not be a primary determinant of individuals' consumption levels. Frequently a concept of equal treat-

ment is the motivating criterion. Thus we hear, "This is an energy shortage beyond everyone's control; those who cannot afford higher fuel prices should not be forced out of their homes as would happen if we allowed the market to operate; we must ration the available fuel evenly," or "We cannot allow scarce federally supported summer jobs to be given out as patronage; in the future we should employ a lottery."²⁵ At other times favoritism, explicit or implicit, may be part of the process. War veterans receive priority in job assignments.

Whether we welcome moves away from traditional competitive allocation schemes, deplore them, or take an intermediate view, we must at least recognize the significant and growing importance of nonmarket allocation schemes in contemporary society. All of the traditional arguments for efficiency are maintained when money is removed from the picture. Indeed, they are reinforced. Arguments that suggest that the pursuit of efficiency—whereby income determines allocations—is likely to result in the trampling of equity are no longer valid. Equity and efficiency are independent considerations, not competitive ones.

Some hopeful results are reported here. Efficiency, defined by the external context to be within the set of outcomes not requiring side payments, can be achieved without using money as a facilitating mechanism. The required allocation procedure is intuitively comprehensible and mechanically workable. As a replacement for a number of existing nonmarket allocation procedures, it could improve prospects for all participants.

Appendix

Summary of Appendices

The paper with full appendices treating the more technical aspects of the analysis is available as Discussion Paper 51, Kennedy School of Government, Harvard University, Cambridge, Massachusetts 02138.

Appendix A.—Proof of the existence of market-clearing prices. The proof makes direct use of Kakutani's fixed-point theorem (see Debreu 1959, p. 26). For an arbitrary nonnegative price vector, an individual's demand will not be an upper semicontinuous function of price because of the constraint that the purchased lottery shares must sum to 1. If we restrict ourselves to price vectors in which at least one component is 0, upper semicontinuity can be proved. The existence of market-clearing prices then follows from the fixed-point theorem, and it follows that the prices can be chosen such that at least one of them is 0. (Special care must be taken to guarantee that the function, a fixed-point of which solves the problem, is defined on a convex and compact set.)

²⁵ After a scandal in Boston in the summer of 1977, the lottery proposal has received serious consideration there.

Appendix B.—Computation of market-clearing prices and allocations. The problem is discussed and difficulties with various direct approaches are pointed out. The computation of individual and aggregate demand functions is described. (No particular difficulties occur here.) Thereafter, we turn to the application of Scarf's (1973) algorithm for computing market-clearing prices. The working of the algorithm is outlined. Then we discuss the sense in which the algorithm solves our problem; that is, in which respects the produced results always will, and in which respects they need not, be close to true market-clearing prices and allocations. The outcome will always be arbitrarily close to the Pareto frontier.

Appendix C.—Conduct of the lottery. The lottery procedure is described formally, and it is proved that it always works. (In particular, it is proved that one can always find a "cycle" on which an adjustment step can be performed.) The procedure does not tell us how different individuals' final assignments are interrelated (they cannot be stochastically independent), which may depend on the arbitrary choice of one among several possible cycles at each adjustment step. It is possible to formulate a limited concept of independence among individuals' assignments; this concept is discussed in the latter part of the Appendix.

Appendix D.—Extensions and generalizations. Straightforward generalizations are described for the cases in which the total capacity of the positions exceeds or falls short of the number of individuals (then, of course, slots will be left vacant or persons unassigned). Somewhat more complicated, but still possible to handle, is the case where there is a limited flexibility in the capacities of the positions. An important generalization allows multiple assignments. If there is no requirement that a person's various assignments be different and if preferences are additive over received positions, our method can be applied. Somewhat more general preferences can also be handled, namely, the case in which a person's preferences for an additional piece of a good depend on how much the person already has of that good. The method cannot, however, be immediately extended to arbitrary preferences.

References

- Birkhoff, Garrett. "Three Observations on Linear Algebra." *Rev. Univ. Nac. Tucuman*, ser. A, 5 (1946): 147–51.
- Dantzig, George. *Linear Programming and Extensions*. Princeton, N.J.: Princeton Univ. Press, 1963.
- Debreu, Gerard. *Theory of Value*. New York: Wiley, 1959.
- Gale, David, and Shapley, Lloyd S. "College Admissions and the Stability of Marriage." *American Math. Monthly* 69 (January 1962): 9–15.
- Gibbard, Allan. "Straightforwardness of Game Forms with Lotteries as Outcomes." *Econometrica* 46 (1978): 595–614.
- Groves, Theodore, and Ledyard, John. "Optimal Allocation of Public Goods: A Solution to the 'Free Rider' Problem." *Econometrica* 45 (1977): 783–809.
- Luce, R. Duncan, and Raiffa, Howard. *Games and Decisions*. New York: Wiley, 1957.
- Quirk, James, and Saposnik, Rubin. *Introduction to General Equilibrium Theory and Welfare Economics*. New York: McGraw-Hill, 1968.
- Scarf, Herb. *The Computation of Economic Equilibria*. New Haven, Conn.: Yale Univ. Press, 1973.

- Tideman, T. Nicolaus, ed. Collection of papers on the "demand-revealing procedure." *Public Choice*, vol. 29, no. 2 (Spring 1977).
- Tideman, T. Nicolaus, and Tullock, Gordon. "A New and Superior Process for Making Social Choices." *J.P.E.* 84, no. 6 (December 1976): 1145-59.
- von Neumann, John. "A Certain Zero-Sum Two-Person Game Equivalent to the Optimal Assignment Problem." In *Contributions to the Theory of Games*, edited by H. W. Kuhn and A. W. Tucker. Vol. 2. Princeton, N.J.: Princeton Univ. Press, 1953.
- Zeckhauser, Richard. "Voting Systems, Honest Preferences and Pareto Optimality." *American Polit. Sci. Rev.* 67 (1973): 934-46.