Resource Allocation Processes

EDITED BY

KENNETH J. ARROW
James Bryant Conant University Professor, Harvard University
LEONID HURWICZ
Regents' Professor of Economics, University of Minnesota

CAMBRIDGE UNIVERSITY PRESS
Cambridge
London · New York · Melbourne

Published by the Syndics of the Cambridge University Press
The Pitt Building, Trumpington Street, Cambridge CB2 1RP
Bentley House, 200 Euston Road, London NW1 2DB
32 East 57th Street, New York, NY 10022, USA
296 Beaconsfield Parade, Middle Park, Melbourne 3206, Australia

First published 1977
Reprinted 1978

Library of Congress Cataloging in Publication Data

Arrow, Kenneth Joseph, 1921-
Studies in resource allocation processes.

Some of the studies prepared in cooperation with other authors.

Includes index.
1. Economics, Mathematical - Addresses, essays, lectures.
2. Mathematical optimization - Addresses, essays, lectures.
I. Hurwicz, Leonid, joint author. II. Title. III. Title: Resource allocation processes.
HB135.A79 330'.01'51 76-9171

ISBN 0 521 21522 6

Transferred to digital printing 2003

CONTENTS

The design of resource allocation mechanisms

LEONID HURWICZ

Traditionally, economic analysis treats the economic system as one of the givens. The term "design" in the title is meant to stress that the structure of the economic system is to be regarded as an unknown. An unknown in what problem? Typically, that of finding a system that would be, in a sense to be specified, superior to the existing one. The idea of searching for a better system is at least as ancient as Plato's *Republic*, but it is only recently that tools have become available for a systematic, analytical approach to such search procedures. This new approach refuses to accept the institutional status quo of a particular time and place as the only legitimate object of interest and yet recognizes constraints that disqualify naive utopias.

A wealth of ideas, originating in disciplines as diverse as computer theory, public administration, games, and control sciences, has, in my view, opened up an exciting new frontier of economic analysis. It is the purpose of this paper to survey some of the accomplishments and to consider outstanding unsolved problems and desirable directions for future efforts.

It is not by accident that the terms "analytical" and "institutional" were only a few words apart in the preceding statement of scientific goals of our inquiry. In the past, especially in the nineteenth century, cleavage developed between analysts who tended to focus on the competitive and monopolistic market models and institutionalists who, either as historians or as reformers, felt the need for a broader framework, but found the existing analytical tools inadequate for their purposes. It is perhaps symbolic that a lecture named after "the father of institutional economics in the United States" [66] should provide a forum for a step toward synthesis of the two approaches.

I should make clear that I do not regard Richard T. Ely as a hundred percent kindred spirit. One reason for this may be seen from the following quotation of his views [17] on mathematical economics:

> No mention has been made of the younger "mathematical school" of political economists, of whom the chief representatives are Stanley Jevons . . . and Léon Walras . . . , because it is difficult to see in their mathematico-economical works anything more than a not very successful attempt to develop further the older abstract political economy. Any advance of the science due to the mathematical character of their method has certainly not yet become widely known, and the writer is much inclined to believe that the works which have advocated the application of mathematics to economics

> form no essential part of the development of economic literature. Certain unreal conceptions and a few definitions are used as bases for mathematical deductions. [p. 60]

Yet I find much to agree with Ely in his broader scientific objectives. However, I shall not go so far as to propose reinstatement into the bylaws of the American Economic Association the platform provisions which he proposed [18] when the Association was being formed:

> 1. We regard the state as an educational and ethical agency whose positive aid is an indispensable condition of human progress. While we recognize the necessity of individual initiative in industrial life, we hold that the doctrine of *laissez-faire* is unsafe in politics and unsound in morals; and that it suggests an inadequate explanation of the relations between the state and the citizens. [pp. 6–7]

What I do sympathize with in Ely's attitude is the desire to view the economic system as a variable and to go beyond analytical frameworks that were unable to cope with this problem. A sharp statement illustrative of the "activist" point of view of that era (quoted by Ely who, however, characterizes it as too narrow) is the following definition, due to the Belgian Emile de Laveleye, dating from 1882: "Political economy may . . . be defined as the science which determines what laws men ought to adopt in order that they may, with the least possible exertion, procure the greatest abundance of things useful for the satisfaction of their wants, may distribute them justly and consume them rationally." I do feel that Ely underestimated the potential of development of theory (and of mathematical theory in particular) to help in this endeavor; but given the lag of a better part of a century in this development, perhaps he can be forgiven.

In what follows I want to focus on developments that are relatively recent, primarily those of the last two decades, and characterized by at least an attempt at rigorous mathematical formulation. First, however, I want to acknowledge the value of work that preceded the recent period.

In spirit, I regard the Utopians, and Utopian socialists in particular, as the initiators of what one might call an "activist" (as well as critical) attitude toward the social system in general, and the economic system in particular. They were, in a sense, the first systems designers in the social sphere. Marx, Engels, and their followers broke with the Utopian socialists. An unfortunate byproduct was the neglect of problems of resource allocation in the ("historically inevitable") socialist economy of the future, with Kautsky something of an exception. In the late nineteenth century there were, however, nonsocialist (and even anti-socialist) economists who tackled the problem in a remarkably objective spirit, among them Pareto, Boehm-Bawerk, and von Wieser.[1] Barone's now famous 1908 paper [9] was at least partly stimulated by Pareto's earlier analysis.

A "second round" of discussion was largely provoked by von Mises' skepticism as to even a theoretical feasibility of rational allocation under socialism.

Oskar Lange's contribution to the debate in the 1930's is well known, but there was a remarkable earlier reply by Jacob Marschak in 1924 [45].

While Lange's line was to be that socialism is as capable of playing the perfectly competitive game as is capitalism, Marschak took the opposite view: capitalism is a world not of perfect competition but of monopolies and cartels, which (in a Schumpeterian spirit) has its good points, especially in the realm of dynamics. Marschak expected similar phenomena under the brand of democratic "socialism" he had in mind and was not depressed by the prospect. He felt that the advantages of imperfect competition would carry over into collectivism. The real issues would be not rational economic calculation, but motivation, stimulation of initiative, and intensity of effort – under an egalitarian system where managers would be democratically elected. (Still, he regarded these problems as less severe than those of "centralistic" socialism.)

In the 1930's, two major lines of development are relevant. One line was the work of Lange, Lerner, and others on resource allocation in a socialist economy, the slightly earlier (1929) paper by F. M. Taylor on trial and error methods, Hotelling's contribution on marginal cost pricing and consumer-producer surplus, and the "new welfare economics" of Hicks, Kaldor, Scitovsky, and others. The other line of development, started in the 1940's, was the mathematization of "classical welfare economics" by Lange [37], Allais [1, 2], Debreu [14, 15], Arrow [4], and Koopmans [34], with the Arrow-Hahn book [5] a recent entry in this series.

I have so far been stressing the ideas oriented toward redesigning the economy of a nation or similar collectivity. But with the enormous growth of private enterprises and governmental bodies, similar issues arise in determining the relationships between the headquarters of a firm and its divisions, or a ministry (department) and its components. Most of the proposed mechanisms are highly relevant in such circumstances, but the team theory model may be particularly appropriate.

There is also a close relationship with information theory and with problems of administrative organization. For linking up resource allocation with information processing and organization, major credit must go to J. Marschak's [46] development of economic theory of information and to Herbert Simon [63] for his work on organization and economic behavior.

Also one should not forget a pioneering effort toward an abstract formalization undertaken by J. B. Kruskal, Jr., and Allen Newell in "A Model for Organization Theory" [36], circulated at Rand Corporation in 1950 but, I believe, never published.

It has been said, only half in jest, that the theory of organization is a field rich in definitions, but short on results in the form of theorems. This is no longer true. There are two categories of results. On the one hand, quite a few specific allocation mechanisms have been invented and their properties, such as feasibility, optimality, and convergence, rigorously established. On the other hand, there are also some more general results, dealing with the possibility or impossibility of various types of decentralized mechanisms, depending on the environments with which they must cope. We shall mention a few of these results, following the sampling of the specific mechanisms and the discussion of a framework required for formulating the more general questions.

I. Specific mechanisms whose properties have been investigated

As promised, we shall now sample, although very incompletely, some of the wealth of specific mechanisms that have been formulated in a rigorous way, mostly during the last two decades - but without forgetting the crucial influences of their less formal predecessors.

We shall confine ourselves to procedures that have been formulated with sufficient precision to avoid ambiguity as to which economic agent says what to whom and when. This makes it possible to determine the informational requirements, as well as convergence and optimality properties.

A major impetus was given to the design of such mechanisms by these developments of the 1940's:

1) activity analysis and linear programming (including the simplex method) - Dantzig, Kantorovitch, Koopmans;

2) game theory, including the iterative solution procedures - von Neumann and Morgenstern, George Brown, Julia Robinson;

3) discoveries concerning the relationships connecting programming (linear or nonlinear), two-person zero sum games, and the long known Lagrange multipliers - Gale, Kuhn, Tucker.

While in economics one deals with goal conflicts due to multiplicity of consumers, linear and nonlinear programming models usually presuppose a single well-defined objective function to be, say, maximized, i.e., a situation corresponding to an economy with a single consumer. So it is not surprising that the mechanisms designed under the influence of programming theory dealt to a large extent with one-objective-function problems and thus failed to face the crucial issue of goal conflict. Nevertheless, one should not underrate their usefulness as a necessary step on the road to the harder multi-objective problem, since the difficulties of the simpler situations do not disappear when goal conflicts are introduced.

We can distinguish two strands here: one, a rather close relative of the programming approach; the other, "team theory" [19-21, 46, 47, 55], more closely related to the theory of statistical inference and decision making. We shall concentrate on the former.

We thus have a situation where there is only one consumer (individual, firm, or even nation) and, hence, only a single utility function to be maximized, but a multiplicity of producers and resource holders. The technological relations (production functions) and limits of resource availability constitute the constraints subject to which maximization must be carried out.

Two difficulties make the problem nontrivial: calculation and information transfer. First, consider the calculation of the maximizing values for the variables of the problem. Assuming even that all the relevant information concerning the parameters of the problem is in the hands of a computing agency, this agency needs a well-defined computational procedure (algorithm) to find solutions.

For linear economies, the simplex method is such an algorithm. For "smooth" (nonlinear) unconstrained maximization problems resembling the task of groping one's way in the dark to the top of a hill, the obvious idea of moving uphill is embodied in the notion of gradient (or "steepest ascent") processes. Equally evident is the fact that a valley between the spot one is at and the peak of the hill would cause trouble; hence, the success of gradient procedures depends on the curvature characteristics of the terrain, a natural requirement being that the hill be dome-like (technically, a strictly concave function).

A natural extension of the gradient process idea was suggested by the famous Kuhn-Tucker theorem associating with a constrained maximum a saddle point, i.e., a maximum–minimum point of the so-called Lagrangean expression; thus the search for a maximum with respect to certain decision variables was converted into a mixture of maximizing and minimizing tasks to each of which one could apply the gradient idea (groping upward in terms of the decision variables and downward with respect to certain auxiliary variables - Lagrange multipliers to the mathematician, shadow prices to the economist). *Statically*, the Kuhn-Tucker result required little more than the concavity (not necessarily strict) of the relevant functions and, hence, was applicable to linear problems as well as the typical smoothly curved pictures of classical economics (Pareto, Hicks). But we shall see that dynamics was more troublesome.

Even when there is an algorithm suitable for calculations by an agency to whom all the data are available, it may be that the information-processing capacity of any one such agency is inadequate because of the size of the problem (the number and complexity of constraints, objectives, and variables). If there are several potential information-processing agencies (and here every human brain qualifies to some extent), we may be saved by devising computing procedures which parcel out the work among them; these are called decomposition algorithms. (In recent years, related problems have been studied in connection with the design of electronic computer utilization under the label of multiprocessing.)

It should be recalled that one of Hayek's [22, p. 212] chief points in summing up the state of the debate concerning the feasibility of a centralized socialist solution was that the number of variables and equations would be "at least in the hundreds of thousands" and the required equation solving "a task which, with any of the means known at present, could not be carried out in a lifetime. Any yet these decisions would . . . have to be made continuously . . . ."

The market-simulation procedure developed by Lange and Lerner may be viewed as an early example of a decomposition algorithm.

From the point of view of the economics of information processing, it is clear that a parceling out of the task may be advantageous even if single agency capacity constraints have not been reached; this may well lower the resource cost and cut down the time required for the completion of the computing process.

But another informational consideration, stressed by Hayek [22, 23], has gained special prominence: the difficulty of placing all the relevant information in the hands of a single agency because information is dispersed throughout the economy. A natural assumption is that, initially, each economic unit has information about itself only: consumers about their respective preferences, producers about their technologies, and resource holders about the resources. An attempt to transfer all this information to a single agency before it starts its calculations is regarded as either impossible (in the sense that much information would be lost) or too costly in relation to the existing accuracy requirements. (One reason for this difficulty is that even the individual units have the required information only in potential form, except for situations corresponding to their past experience: e.g., firms know only certain parts of their production functions. It is easier to use "localized" procedures which require an exploration only of the relevant parts of the individual units' maps; but such localization is impossible if whole maps are to be conveyed to the single computing agency at the beginning of the computing process.)

If the economic units which initially are the only ones with information about themselves are also capable of carrying out calculations, it is natural to seek computing procedures that would both minimize the need for information transfers and also parcel out the tasks of calculation. This is what informationally decentralized procedures are meant to accomplish.

In the 1930's it would have been most natural to start with "smooth" strictly concave economies (diminishing marginal utilities and returns), without kinks or corner solutions. But around 1950, linear models were in fashion. Furthermore, the simplex method was available and proved to be convergent. Since the simplex method, applied to the economy as a whole, lacked informational decentralization, a search for an alternative was bound to occur. For the economist, an obvious candidate was a simulated (perfectly competitive) market process à la Lange, Lerner, and (specifically in the context of a linear economy) the Koopmans model with a helmsman representing the consumer, production managers maximizing profits, and resource custodians adjusting prices according to excess demand.

From a static point of view, the equilibrium of such a process would be optimal. But if the initially proposed process and quantities were "wrong," would there be convergence to equilibrium? To make this question meaningful, one must specify the dynamics of the adjustment process, e.g., by how much prices are to be raised per unit of time given the magnitude of excess demand, etc. A pioneering model of this type is due to Samuelson [58] who postulated a system of differential equations in which prices vary proportionately to excess demand, and resource use rises when low resource prices yield positive profits. He immediately noted that this dynamic system would behave like a frictionless pendulum, i.e., would not converge to an equilibrium position. Whether we are thinking of computations or of designing an economy, we must look further.[2]

Samuelson's discovery posed a challenge: can an informationally decentralized convergent allocation process be designed for linear economies? (Of course, there is also the problem of designing such processes for economies with increasing returns. But here the difficulties are bound to be serious, since the competitive mechanism lacks even the usual static properties.)

One line of attack involves the replacement of the fixed (that is, parametric) price idea by that of a price *schedule* and is applicable not only in the linear (constant returns) case, but also under increasing returns. (See the modified Lagrangean Arrow-Hurwicz approach below.) Another approach, to be discussed first, retains the parametric prices; it grew out of linear programming techniques, with the Dantzig-Wolfe decomposition method [13] its earliest example.

The Dantzig-Wolfe economy has special features which provide scope for the decomposition approach. These are the usual features of a resource allocation model without technological externalities, in which the objective function is a sum of the contributions of the individual units and in which certain resources must be utilized by all units. (Both the objective function and the overall constraints are "additively separable.") The mechanism may be viewed as a dialogue between the producing units (who know their technologies and contributions to the objective function) and a "center" which knows the total resources available. One aspect of the dialogue is that the center proposes tentative resource prices and the producing units develop corresponding profit-maximizing production programs (with prices treated parametrically). In the light of these programs, the center revises the proposed prices. Because of the linear character of the economy, both the center and the producing units can use linear-programming (primal and dual) techniques, and an equilibrium is reached in a finite number of steps. So far, we may regard the algorithm as a variety of the market (parametric price) process. But there is a difference. The final allocation will not necessarily correspond to the final production programs of the producing units. Rather, the center will "order" each producing unit to undertake a program which "mixes" (averages) the final proposal with several previous ones.[3]

As pointed out by Baumol and Fabian [10], the procedure can be extended to situations where constraints pertaining to single producing units are nonlinear, while the overall constraints pertaining to resources needed by all units remain linear. In this case, however, the units must have a computational algorithm for their nonlinear problem since the simplex method can no longer be used. For other nonlinear economies, and especially those with increasing returns, different processes had to be sought.

Before we look at those, however, let us examine a mechanism designed specifically to guide a *linear* economy but in a manner that partly reverses the roles played by the "center" and the "periphery" (the producing units), the process due to Kornai and Lipták [35]. The assumption concerning the economy, as in the Dantzig-Wolfe model, is that of "block angularity," i.e., there are subsets of constraints each pertaining to a given sector and also resource constraints affecting the whole economy. In the dialogue, the center proposes allotments of scarce resources to the various sectors; then each sector responds with shadow prices (marginal rates of substitution) minimizing the value of the allotment subject to sectoral dual constraints (nonprofit condition for every sectoral activity). The center's aim, on the other hand, is to maximize the contributions of the sectors to the objective function, i.e., to maximize the value of the allocated resources at the shadow prices received from the sectors, subject to the limitation of available resource totals.

Taking advantage of the equivalence of linear-programming programs and games, Kornai and Lipták, by structuring the dialogue as a fictitious game, are able to establish convergence to an equilibrium with any desired degree of accuracy, though (unlike in the Dantzig-Wolfe procedure) without reaching the equilibrium in a finite number of steps. In addition to the latter disadvantage, it has also been pointed out [30] that the Kornai-Lipták procedure is not completely informationally decentralized, since each sector's resource sectoral allotments must be large enough ("evaluable") to assure the existence of a feasible solution for that sector.

One advantage claimed for the Kornai-Lipták procedure is that it may be computationally manageable where alternative decomposition algorithms are not. I regret that I have not had an opportunity to look into this question. But another feature of this process is of great interest to the economist. This is the fact that the center, instead of simulating the market as does the Taylor-Lange-Lerner mechanism, specifies quantitative input and output targets or restrictions, while the sectors supply the center with productivity information in the form of shadow prices. This appears more in line with many observed planning practices and thus may provide a useful descriptive model.

There are several other mechanisms, in general designed for nonlinear economies, which are also of the "quantity-guided" type (as distinct from the "price-guided" type), that is, where the center sends out messages concerning quantities (e.g., targets) and the periphery (the producing units, sectors) responds with marginal entities or shadow prices. Informationally, since the center is sending different quantity messages to different sectors, its total signals are of higher dimensionality than in price-guided systems where the same message (price vector) goes out to all sectors. (We must bear in mind that each quantity vector has the same dimension as the price vector.) Whether this difference is significant is somewhat controversial; the negative has been strongly expressed by Marglin [44] who has constructed quantity-guided (called by him "command") counterparts of certain price-guided processes. One of Marglin's mechanisms requires the center to allocate the scarce inputs on the basis of information obtained from the producing units concerning their marginal productivities and their excess demands. Adjustment ceases when aggregate excess demand is zero and the marginal productivities of producers are equalized. (A similar process was proposed by Heal [24].)

A process, characterized by a mixture of price- and quantity-guided elements and also due to Heal [25], is particularly interesting because (with some qualifications) it converges to optima even for nonconvex economies, in particular for increasing returns. An essential informational feature is that certain functions of each producing unit's marginal productivities (roughly, its shadow prices for particular resources) must be conveyed to the center. The center can then calculate improved resource allotments, or else it may calculate and send to the units a resource price (the same for all units) and so enable them to determine their respective resource requirements. The latter option is, of course, informationally more decentralized: it requires fewer message transfers from the center to the periphery, and fewer computations are carried out at the center. (Marglin had a similar process for a more restricted class of economies. It is not clear whether his process could be adapted to corner maxima in nonconvex cases.)

Heal's process has the further merit that if the initial allocation is feasible, so are all the later ones, thus satisfying a Malinvaud postulate. (The same seems true of Marglin's process, although, unlike Heal, he does not assume the initial position to be feasible.) The maintenance of feasibility is simple in models without intermediate goods because the procedure always allocates all available resources and producers are required to stay on their efficient frontiers. The matter gets more complicated when intermediate goods are introduced and only special cases appear to have been dealt with so far by this approach.

Processes in which the center specifies quantities and the peripheral units convey their individual marginal rates have also been used in models where public goods are among those to be allocated. I shall mention three treatments of this case, due to Drèze and de la Vallée Poussin [16], Malinvaud [43], and Aoki [3]. In the versions known to me each is somewhat specialized: Malinvaud and Drèze-Poussin have only one producer but many consumers, Aoki only one consumer; also Drèze-Poussin deal primarily with the case of only one private good, although they indicate how the results may be generalized to more. Aoki's economy is closest to those we have been considering so far because it has only one consumer (the center), hence, no income distribution problems; it also has many producers. His mechanism uses price-guidance for private goods and quantity-guidance for public goods. A producer develops production plans that maximize net revenue given the central "guidelines" (prices for the private goods, quantities for the public goods), and conveys to the center his demands for private goods and marginal evaluations, including marginal cost, for public goods. The center, in turn, adjusts the price of each private good according to the difference between its marginal utility and price (as in the Arrow-Hurwicz gradient process); the targets for public goods are increased in proportion to the net aggregate of marginal valuations (users' minus producers'); thus the center combines the functions of the helmsman and resource custodian of the Koopmans model for private goods with target setting for public goods. The other two processes have the same adjustment rule for public goods targets, but differ in other respects. In particular, they specify the rules of income distribution.

All three processes converge (at least locally) under suitable convexity assumptions concerning the environment. Somewhat paradoxically, Malinvaud's process does not seem to satisfy his desiderata of feasibility maintenance and monotonicity, while the other two do.

I shall now go back to the price-guided processes, but more briefly because they came earlier and are better known. Here again, I shall focus on the one consumer case. For a linear economy, Koopmans [33] described the functioning of such a mechanism in the spirit of the Taylor-Lange-Lerner rules by setting up an "allocation game" to be played, in an informationally decentralized manner, by a helmsman (setting the prices of final goods and thus representing consumer preferences), commodity custodians (adjusting the prices of resources according to excess demand), and activity managers who determine the production programs. Koopmans' adjustment rule (similar to Samuelson's) is that managers expand profitable activities and curtail those bringing losses; in a constant returns economy this is equivalent to profit maximization. Koopmans stressed that "the dynamic aspects of these rules have on purpose been left vague." We know that Samuelson's experiment in this direction yielded nonconvergent oscillations.

On the other hand, in an economy where all functions (including the utility indicator) are strictly concave (i.e., we have diminishing returns), similar rules produce a process with the desired stability properties. Utilizing the notion of gradient approach to the saddle point of the Lagrangean expression, Arrow and Hurwicz [8] used the following rules: the helmsman, taking the prices of desired commodities as given, changes each final demand at a rate equal to the difference between its marginal utility and price; each manager, again taking prices as given, changes the scale of his process in proportion to its marginal profitability; each commodity custodian varies the price of his commodity in proportion to excess demand. (I am omitting modifications pertaining to corners and zero prices.) A limiting form of such a process is the price-adjustment method in which prices are varied as before, but both the helmsman and each manager reach (as against merely moving toward) the values of their decision variables which maximize their respective objective functions: for the helmsman, the difference between utility and price; for the manager, the level of profit. The familiar Walrasian competitive process is a variant of such price-adjustment in which the demand for final goods is determined by utility maximization subject to a budget constraint with specified income or wealth.

Although the gradient process is informationally decentralized and converges to an optimum under strict concavity assumptions, it has certain disadvantages. To begin with, it is formulated in continuous time, while realistic mechanisms operate more naturally by iterations, i.e., with a discrete time parameter.

But it is possible to construct a discrete time parameter counterpart of the gradient process. In fact, this was done by Uzawa [67] and further elaborated by Malinvaud [42] who, however, pointed out another undesirable feature of the gradient process: although it converges to a feasible solution, its "interim" proposals are, in general, not feasible. In other words, while the process pushes the participants toward compatibility in their claims on resources, they may be demanding either more or less than the total available while the process is going on. Thus if the gradient process were to be interrupted at a finite time, there might be a problem of reconciling incompatible claims.

Malinvaud then formulated a desideratum, viz., that (discrete time parameter) adjustment processes yield feasible solutions after a finite number of iterations. He then proceeded to construct two processes (for different environments) which satisfied this desideratum. In both cases he assumed that the initial proposal, serving as the point of departure for the iterations, was feasible. In effect, he was paying a "price" for the feasibility of all his interim proposals - namely, he assumed that the center had an additional piece of information: a feasible point of departure; the Arrow-Hurwicz gradient process, on the other hand, was designed for situations where this information was not available to the center.

It is worthwhile to become acquainted with Malinvaud's second procedure. The center proposes prices to the producing units which, in turn, determine production plans maximizing the value of the firm's output in terms of those prices. The center then builds up its picture of each unit's production set (see Fig. 1) by taking all convex mixtures of its previous proposed input-output vectors, together with the initial feasible vector, assumed known to the center. (Since the production sets are assumed convex, this yields an increasing subset of the unit's true production set. But, again, there is an informational price to be paid: the center must accumulate, on a disaggregated basis, all past proposals from the units.) Treating its pictures of the production sets as if they were the actual sets, the center then maximizes its utility function subject to the resource availability constraint and proposes a new set of prices corresponding to the relevant marginal rates of substitution.

It could perhaps happen that, even after several iterations, the only production program compatible with resource constraints is the one originally assumed known to the center; or that, even if new feasible programs are generated, their utility is no higher than that of the original known feasible allocation. But it was shown by Malinvaud that, as the number of iterations goes to infinity, the utility associated with the corresponding plans tends to the upper bound of its feasible values.

Roughly speaking, the center constructs plans which would be nearly optimal for the economy if its images of the individual production sets were sufficiently close to correct. The informational price paid is the need for building up these images; in effect, the information concerning the production functions is being transferred to the center, although on an installment plan. This differs from the usual informationally decentralized procedures where the center does not accumulate such information and never knows more than the structure of the production set in the neighborhood of the current proposal.