COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 1068

THE ALLOCATION OF RESOURCES IN THE PRESENCE OF INDIVISIBILITIES

Herbert E. Scarf

January 1994

THE ALLOCATION OF RESOURCES IN THE PRESENCE OF INDIVISIBILITIES[1]

Presented at the Annual Meeting of

the American Economic Association

Boston, Massachusetts, January 3, 1994

Herbert E. Scarf

I'd like to talk to you today about the problems presented to economic theory by the presence of indivisibilities in production. I will go about this in a rather leisurely way, beginning with a discussion of the role played by competitive prices in detecting optimality when production takes place under constant returns to scale; then illustrating the failure of prices to perform this task when indivisibilities are significant; and, finally, suggesting the replacement of the pricing test by a specific quantity test. It is my hope that continued study of these quantity tests will increase our understanding of the division of labor in a large firm.

Let me begin this discussion with a simple observation: If the economic environment in which we find ourselves were not continually changing, there would be no need for our marvelous science of economics. In a world in which there are no changes in population size, or in the availability of raw materials and other factors of production -- in which privately held technical knowledge is not disseminated to potential competitors; in which no innovations in productive techniques arise; in which there are no new products, no changes in consumer preferences and no earthquakes, floods or other natural catastrophes -- in such a world optimal behavior, for every one of the agents in the economy, is simply to continue with precisely the same routine, habitual behavior engaged in yesterday. It doesn't matter whether our daily agenda is the result of a Stalinist central plan, or a solution of a fully specified Walrasian model of equilibrium on a super computer, or inspired by the institution of perfectly competitive markets; it doesn't matter because the absence of novelty implies that our economic activities, though possibly complex, are essentially repetitive.

---

[1] I would like to thank Truman Bewley, William Brainard, Alvin Klevorick, William Nordhaus, T.N. Srinivasan and James Tobin for their thoughtful comments.

But the world around us isn't static, and changes are constantly presenting themselves. Estimating the consequences and the merits of changes in our economic environment is the bread and butter of economic theory; we do it all the time. We do comparative statics at the level of the firm, when we estimate the consequences of a change in factor endowments or in the price of a valuable input into production. We engage in comparative statics and dynamics for the economy as a whole when we examine the consequences of the dramatic increase in the price of imported oil in the latter part of 1973, or the second oil shock following the fall of the Shah, or the dismantling of AT&T, or a massive change in income taxation within the United States, or the NAFTA. If these consequences spread throughout the entire economy, we evaluate them by assessing their effects on the well-being of the members of the community.

We have a remarkable paradigm for assessing well-being that has been passed on to us by generations of economic theorists and utilitarian philosophers. The utilitarian calculus, in its modern ordinal version, provides us with a simple test for evaluating the merits of a proposed change in economic activities: The change should be accepted if it has as an immediate consequence -- or one that can be brought about by a suitable redistribution of income -- an increase in the well-being or utility of all of the members of society.

It may be fashionable, at present, to dismiss the utilitarian test for the evaluation of a novel project on several grounds. The test requires the possibility of major income redistributions that may not be politically viable -- the movement of a clothing manufacturer from a Northern mill town to a lower wage region of the South may result in a potential Pareto improvement, but I know of no instances of an appropriate compensation to those employees whose jobs have been lost. Job retraining is a poor substitute for the income redistributions necessary to retain previous levels of utility. And there are serious problems about maintaining effective incentives if lump sum transfers of income are made independently of effort and the supply of productive factors.

In spite of these and other doubts, I personally consider the welfare test to be an extraordinary intellectual construction -- one which permits us to focus our discussions about the potential merits of a novel economic proposal. Last summer, for example, I participated in an extended discussion with a very distinguished high energy physicist about the Super Conducting Super Collider. In the middle of our conversation, it became quite clear to me that the community of physicists in favor of the project had been unable to establish any ground rules about what constituted a compelling argument for the project. It wasn't the case that they had no arguments in favor of the Collider; they had many of them. But there was no prior agreement or understanding between the proponents of the Super Collider and their audiences about how to evaluate the merits of any particular argument. We were, of course, talking about a project with a set-up cost of ten or more billion dollars and with a high yearly budget, a project which presumably could have been postponed for several years, or possibly for decades. Any justification for this project should have drawn on mutually agreed lines of argumentation that could conceivably have accepted this project and rejected one whose costs were orders of magnitude larger.

Our profession does have such a line of discourse. It may, admittedly, be difficult to carry out the welfare test in an instance as complex as the Super Collider; the Collider is, after all, a public rather than a private good, and it is, moreover, one whose potential benefits are extremely difficult to predict. But if we are to use arguments other than a direct appeal to the politics of job creation, some variant of the utilitarian calculus must be in the background of the discussion to differentiate this project from alternative uses of public funds.

The utilitarian test is much easier to carry out when more conventional economic projects are proposed. The test actually leads to a simple exercise in the calculation of profitability, which, in my opinion, is one of the major theorems of microeconomic theory, a theorem which is not entirely obvious to the man on the street or even to professional economists.

Suppose that we are contemplating a hypothetical economic situation which is in equilibrium in the purest Walrasian sense. The production possibility set exhibits constant returns to scale so that there is a profit of zero at the equilibrium prices. Each consumer evaluates his income (or wealth) at these prices and market demand functions are obtained by the aggregation of individual utility maximizing demands. The system is in equilibrium in the sense that demand equals supply for each of the goods and services in the economy.

Suppose that a technical advance is made resulting in the discovery of a new manufacturing activity -- one which produces a good whose price is already known, at a new location, with different materials, with less expensive labor or with more sophisticated machinery. Shall the new activity be used?

The word "shall" in this question is the same word as in the question, "Shall the Super Conducting Super Collider be built?", and the utilitarian test can be applied by inquiring whether the new activity can be combined with a plan of income redistributions in such a way as to make all consumers better off than they had previously been. On the face of it, this sounds as if we must solve a complex mathematical programming problem; but, in fact, the question has a remarkably simple answer: <u>If the activity is profitable at the old equilibrium prices, then there is a way to use the activity at a positive level so that with suitable income redistributions, the welfare of every member of society will increase</u>. There is no necessity to determine the new equilibrium prices arising after the activity is introduced; the current prices will do. And conversely, <u>if the new activity makes a negative profit at the old equilibrium prices, then there is no way in which it can be used to improve the utility of all consumers, even allowing the most extraordinary schemes for income redistribution</u>. This is an astonishing mathematical theorem, which I urge you to try to prove. I often ask about this theorem in graduate oral exams in microeconomics. The second assertion takes about two lines of proof; the first is more subtle: three lines of proof and a figure will do. I have never seen this theorem, which seems to me to be one of the important theoretical arguments in favor of private enterprise, in any textbook on economics.

It may be worth remarking that if seventeen new activities are presented simultaneously, the pricing test can be applied to the collection of activities in an arbitrary sequence, without regard to decisions made about the remaining activities. If none of the seventeen activities makes a positive profit at the old equilibrium prices, then no subset of them can be used, along with income redistributions, so as to improve everyone's economic lot. If any one of the activities makes a positive profit, then some welfare improvement is surely possible. The activity can be introduced, a new equilibrium determined -- with Pareto improving income redistributions -- and the pricing test can be applied to the remaining activities. This is an extraordinarily decentralized test when one realizes that it presumably could be applied to every minor innovation on the shop floor of a large firm simply by evaluating its profitability in terms of prevailing market prices.

The market test sounds very much like a step in the simplex method for solving linear programs. An activity analysis model of the economy or a firm is given, along with a specified factor endowment and an objective function which is to be maximized subject to the constraint that the factor endowment is not exceeded. In a linear programming problem, a feasible solution to the constraints is proposed, and prices are found yielding a zero profit for the activities in use. The current solution is optimal if and only if the remaining activities make a profit less than or equal to zero.

The simplex method is an extremely efficient algorithm for solving linear programs: Programs involving thousands of variables can be routinely solved on a personal computer by high school students. But what is even more significant for us as economists is that this effective computational procedure is based on an evaluation of profitability identical to that performed by competitive markets. A visitor from another planet who was taught the simplex method for solving linear maximization problems would inevitably be led to the use of prices and profitability to detect optimality. An algorithm -- a mathematical technique for solving maximization problems -- suggests an institution which is central to the way in which we organize our economic lives.

Is this suggestion of a major institutional structure an accident of the simplex method or can it be expected from other computational procedures as well? Is it a reasonable research strategy to address an area of economic theory which is not fully understood -- at least by me -- to cast it in the form of an optimization problem, and to hope that algorithms for its solution will produce a conceptual framework that is relevant to the original economic problem? I'm not sure, but it is a strategy that I have followed for a number of years in an attempt to increase my understanding of the problems posed for economic theory by indivisibilities and economies of scale.

Both linear programming and the Walrasian model of equilibrium make the fundamental assumption that the production possibility set displays constant or decreasing returns to scale; that there are no economies associated with production at a high scale. I find this an absurd assumption, which is contradicted by the most casual of observations. Consider the following parody: Taken literally, the assumption of constant returns to scale in production implies that if technical knowledge were universally available, we could all trade only in factors of production, and assemble in our own back yards all of the manufactured goods whose services we would like to consume. If I want an automobile at a specified future date, I would purchase steel, glass, rubber, electrical wiring and tools, hire labor of a variety of skills on a part-time basis, and simply make the automobile myself. I would grow my own food, cut and sew my own clothing, build my own computer chips and assemble and disassemble my own international communication system whenever I need to make a telephone call, without any loss of efficiency. Notwithstanding the analysis offered by Adam Smith more than two centuries ago, I would manufacture pins as I needed them.

If production really does obey constant returns to scale, there is nothing to be gained by organizing economic activity in large, durable and complex units; in short, there is no economic justification for the existence of firms. Competitive markets would set prices for manufactured goods at every stage of production and cash would be exchanged, or accounts would be reckoned, as goods moved from one task to another. Every step in a complex manufacturing process would be tested for profitability by itself, without regard to its relationship to other potential improvements.

I am, I believe, not alone in thinking that the essence of economies of scale in production is the presence of large and significant indivisibilities in production. What I have in mind are assembly lines, bridges, transportation and communication networks, giant presses and complex manufacturing plants, which are available in specific discrete sizes, and whose economic usefulness manifests itself only when the scale of operation is large. If the technology giving rise to a large firm is based on indivisibilities, then this technology can be described by, say, an activity analysis model in which some of the activity levels are required to assume integral values, 0,1,2,..., only. When factor levels are specified and a particular objective function is chosen, we are led directly to that class of difficult optimization problems known as integer programs.

For a theorist, the major problem presented by indivisibilities in production is the failure of the pricing test for optimality or for welfare improvements. Let us return to our previous discussion of the economy which is in full Walrasian equilibrium, and imagine, as before, that a new activity is discovered. But let us now assume, in contrast to our earlier example, that this new activity can only be run at an integral level. One can argue easily that if the activity makes a negative profit at the old equilibrium prices, then there is no way to use it at a discrete or continuous level so as to improve the utilities of every agent in the economy. The problem arises with the converse; it is perfectly possible that the activity make a positive profit at the old prices and still not be capable of being used at a discrete level to yield a Pareto improvement. And even more problems arise if seventeen activities are presented to us, all of which must be run at an integral level. A welfare improvement will typically require the selection of a subset of the activities, some of which are profitable at the old equilibrium prices, and some of which are not. There is no algorithm based on prices and profitability which permits us to make a sequence of welfare improvements, by introducing one activity at a time, or even to detect which activities should ultimately be used. There is no pricing test in the presence of indivisibilities.

The absence of a pricing test is a truth that must be confronted. It certainly does imply that total decentralization by means of competitive prices is impossible if the technology involves serious indivisibilities. In my own view this is a compelling reason for the existence of large firms, and what it suggests to me is that some serious insight about the large firm might be gained by considering such a firm to be essentially an algorithm for the solution of mathematical programming problems in which some of the variables are restricted to integer values. I should hope that insights from this source would complement other insights about the functioning of large enterprises that are presented in a narrative rather than mathematical form, that are based on a careful analysis of particular historical cases, or that involve flows of information in hierarchical structures. The subject is sufficiently complex so that many voices should be heard.

At this point, it may be useful to look at a numerical example. Let us consider a problem involving a single good that can be produced by a variety of technologies. Each technology is embodied in a particular type of manufacturing plant with a specific cost of construction, with a specific capacity and with a specific unit cost of manufacturing. The level of demand for the product is given exogenously, and we are required to construct a series of plants and to manufacture sufficient product to satisfy this demand at minimum cost. We now have a mathematical programming problem in which some of the variables, the number of plants of each type to be constructed, are integral, and the remaining variables, the amounts manufactured at each plant, are continuous.

The example is artificial in many ways. Perhaps its most serious flaw is the obvious lack of any dynamic considerations. The construction cost is presumably paid at the time of construction, when a capacity for producing the maximum output per period is established. But demand for output manifests itself in a sequence of periods over time, possibly in a predictable though varying fashion, or possibly with a good deal of uncertainty. Moreover, it is plausible to assume that manufactured goods can be kept in inventory, at some cost, so as to satisfy future demand. These elements can certainly be introduced into our problem, but with a considerable increase in complexity. In order to make my points as simply as possible, I will assume that demand is constant over time and that no inventories are kept; the unit costs may then be thought of as the discounted sum of unit costs incurred over time as this constant demand is satisfied.

For fixed construction costs, capacities and unit costs, the optimal construction plan depends crucially on the level of demand. Some levels will call for considerable excess capacity in various plants, and other levels will not. How can we tell whether a proposed construction and manufacturing plan, which meets the demand requirement, does, in fact, minimize total cost? Competitive prices will not work for this class of problems. There is only one option: the price test must be supplemented, or replaced, by an effective quantity test.

At this point, I have an expository difficulty about which I must be quite explicit. I would like to present an elementary example illustrating the particular quantity test required to demonstrate optimality without being cluttered by too much detail; this naturally leads to an example with a small number, say, two types of plants. But programming problems with only two integer variables are quite easy to solve. In particular, when there are only two types of plants, the saving in cost achieved by a truly optimal solution is quite small compared to the cost of approximately optimal solutions, which are themselves quite easy to find. This is not true for larger problems, and I ask your indulgence on this issue.

With this caveat in mind, let us consider an example involving only two types of plants. The first type of plant -- the Smokestack plant -- is of ancient design, huge, made of red brick with steam pouring from its chimneys; it has a large capacity, is moderately inexpensive to construct per unit of capacity and has a fairly high marginal cost of production. The second plant -- the High Tech plant -- is a gleaming marvel of computerized technology; it has a capacity of medium size, is expensive to set up per unit of capacity, but has a lower marginal cost of production.

	Smokestack	High Tech
Capacity	16	7
Construction Cost	53	30
Marginal Cost	3	2
Average Cost	6.3125	6.2857

If capacity could be built continuously rather than in discrete units, the cost per unit of capacity in the Smokestack plant would be 53/16 and the cost of supplying a unit of demand would be 53/16 + 3 = 6.3125. The average construction and manufacturing cost from a High Tech plant is 30/7 + 2 = 6.2857. What is, of course, uncomfortable about the example, is the closeness of these two average costs.

If plants could be constructed at an arbitrary size, the market test -- using either average or marginal cost as a criterion -- would require that all demand be satisfied from High Tech plants alone. But the optimal solution is considerably different if plants must be built in discrete sizes, and the pricing test for optimality fails dramatically. The following table illustrates the solution for an interval of demand values:

Demand	#Smokestack	#High Tech	Output 1	Output 2	Total Cost
55	3	1	48	7	347
56	0	8	0	56	352
57	1	6	15	42	362
58	1	6	16	42	365
59	2	4	31	28	375

60	2	4	32	28	378
61	3	2	47	14	388
62	3	2	48	14	391
63	0	9	0	63	396
64	4	0	64	0	404
65	1	7	16	49	409
66	2	5	31	35	419
67	2	5	32	35	422
68	3	3	47	21	432
69	3	3	48	21	435
70	0	10	0	70	440

Of course, if the problem were less artificial we would expect a more stable sequence of optimal solutions. In particular, if the capacities at both plants were larger, the number of plants of each type would be considerably less sensitive to the level of demand; a given configuration of underutilized plants would be optimal for a large interval of demands.

It can easily be shown that the number of Smokestack plants becomes a periodic function of demand after some point ($d = 91$ in this example). But it is clear from this table that the optimal integer solution cannot be obtained simply by rounding the optimal solution to the linear programming relaxation.

Let us focus on a particular value of demand, say, 60, for which the optimal solution is to build 2 Smokestack plants, 4 High Tech plants, and manufacture 32 and 28 units respectively, for a cost of \$378. Suppose that an alternative solution had been proposed: that we build 3 Smokestack plants (at a cost of \$159 and providing a capacity of 48), 2 High Tech plants (at a cost of \$60 and capacity of 14), and that we manufacture 46 units at the Smokestack plant and 14 units at the High Tech plant, for a total cost of \$385. Is there a quantity test revealing that this proposal, which does satisfy the demand of 60, is not optimal?

The most elementary quantity test is to plot the point (3,2) in the plane, and examine its 8 neighbors, which are obtained by increasing or decreasing the number of plants of each type by unity. In other words, for any particular feasible construction plan given by a pair (#Smokestack plants, #High Tech plants), we examine those alternative construction plans obtained by adding to this pair of integers each of the 8 vectors:

(1,0),(1,1),(0,1),(-1,1),(-1,0),(-1,-1),(0,-1),(1,-1)

and testing each one of them to see whether it produces another feasible plan at lower cost.

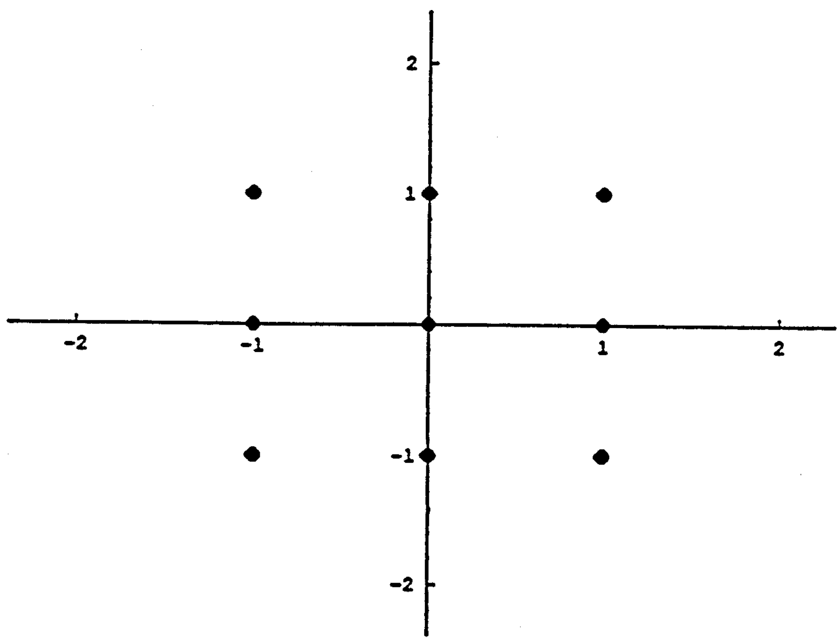

Figure 1

You will notice from Figure 2 that of these 8 points, four of them: (4,2), (4,3), (3,3) and (4,1) are more expensive than (3,2) and the plans associated with the remaining four points (2,2), (2,1), (3,1) and (2,3) do not provide sufficient capacity to satisfy the demand of 60. This local quantity test fails to detect the fact that (3,2) is not an optimal solution when the demand is 60.

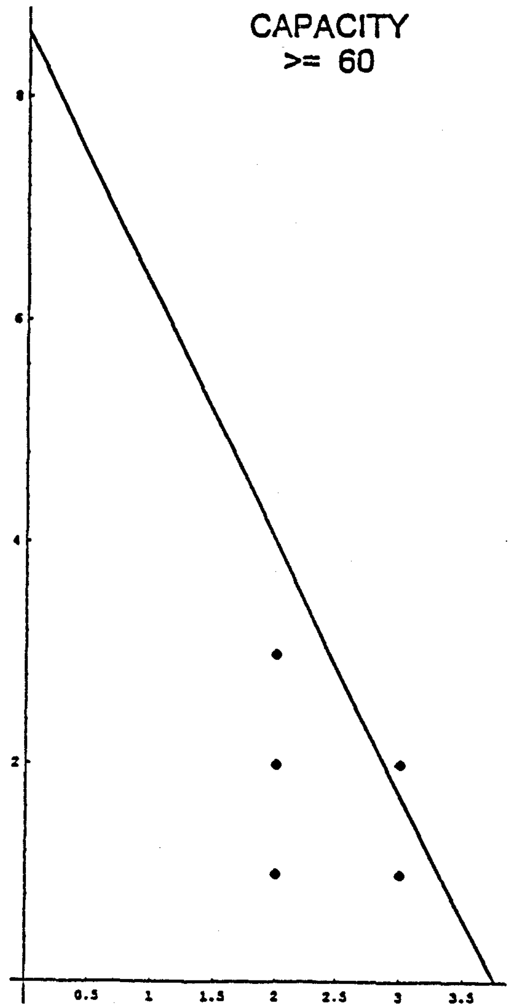

Figure 2

For a quantity test to detect optimality it must be based on an examination of a set of neighbors that are related in some intrinsic fashion to the underlying problem, rather than being merely adjacent in an elementary geometric sense. For our problem, there is a <u>unique, minimal</u> set of neighbors all of which must be examined to be certain about detecting optimality for any level of demand. They are obtained by adding each of the following neighbors, or their negatives, to the proposed plan:

$$(0,1),$$
$$(1,0),(1,-1),(1,-2),$$
$$(-1,3),(-2,5),(-3,7),$$
$$(4,-9).$$

This set of neighbors has the important property that if one of them is not examined, then there will be some level of demand and some feasible solution which is falsely claimed to be optimal. If all of them are examined, the quantity test will yield the optimal solution for any level of demand.

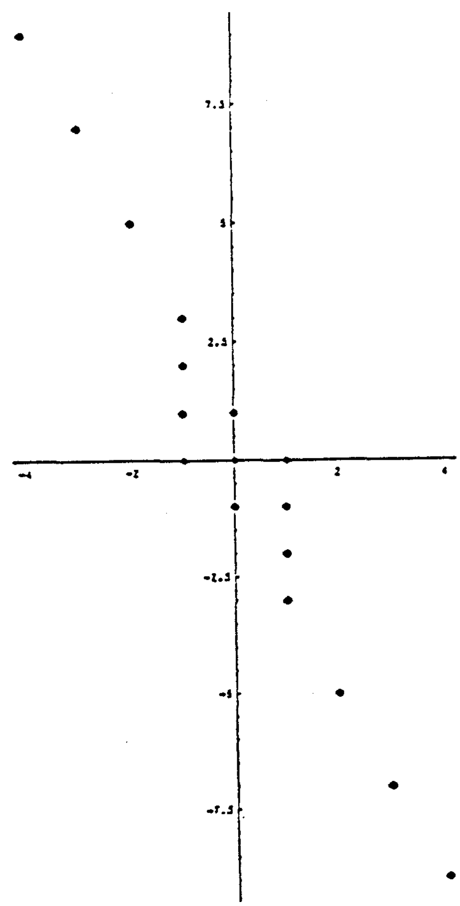

Figure 3

In our example, the non-optimality of the plan (3,2), for a demand of 60, is easily seen in Figure 4 by subtracting the neighbor (1,-2) from (3,2), reaching the new plan (2,4), which is feasible, lowers cost, and, in this instance, is optimal.

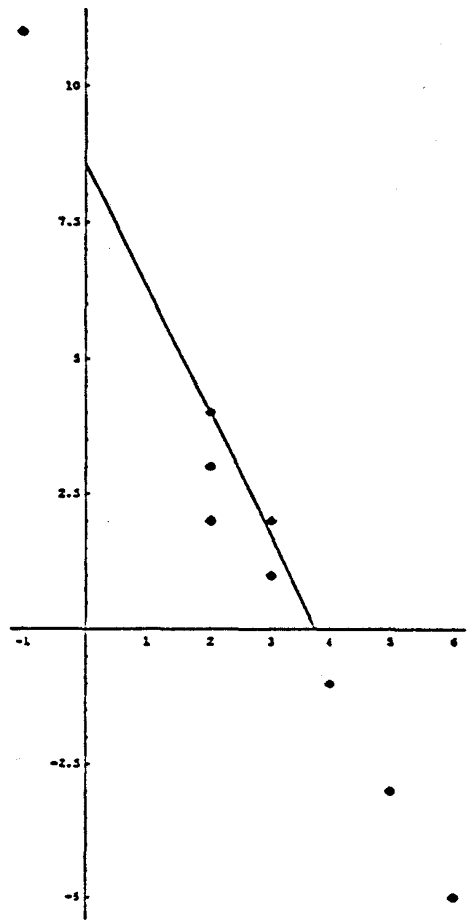

Figure 4

There is a clear algorithm suggested by these considerations: 1) Propose some construction plan which produces a capacity sufficient to meet demand; 2) If one of its neighbors in the unique minimal test set is also feasible and leads to a lower cost, move to that alternative plan; 3) If there are no such neighbors, the original proposal is optimal.

It can be seen that these neighbors are closely related to the discrete analogue of marginal products. As the demand level increases, the optimal construction plan will either be unchanged or move to a new plan which is obtained from the previous plan by adding one of these neighbors.

I don't mean to be unduly mysterious about this concept of neighbors, so let me be more specific about the role that they play in detecting optimality. Consider the neighbor (1,-2). If, as in our example, we subtract this neighbor from a proposed construction plan, we obtain a new plan with 1 less Smokestack plant and 2 more High Tech plants. There will be a net loss in capacity of two units and a net increase in construction costs of \$7. But the 2 additional High Tech plants are capable of manufacturing 14 units at a cost of \$28; these 14 units were previously manufactured at the Smokestack plant for an additional \$1 per unit. It follows that there is a cost saving of \$7 associated with this decrease in capacity of two units.

This "marginal" change would result in a decrease in cost if the original plan had at least 2 units of excess capacity and used at least one Smokestack plant; under these circumstances the change should certainly be adopted and we should move to a new solution with lower cost (as we did in our example in moving from the configuration (3,2) to (2,4)). If we examine the other members of the minimal test set, the decrease in capacity and cost obtained by subtracting each of them from a proposed solution is given in the following table:

Neighbor	Capacity	Cost
(0,1)	7	23
(1,0)	16	53
(1,-1)	9	30
(1,-2)	2	7
(-1,3)	5	16
(-2,5)	3	9
(-3,7)	1	2
(4,-9)	1	5

assuming that all High Tech plants are used to full capacity. The set of neighbors provides us with all of the discrete tradeoffs necessary to verify optimality. Of course, it is not obvious -- nor can I make it obvious without a lengthy, technical argument -- precisely why no other discrete tradeoffs are necessary.

These observations are quite general. Subject to very mild conditions, an arbitrary activity analysis model, with integral activity levels, has associated with it a unique, minimal neighborhood system which depends solely on the technology matrix and not on the specific factor endowment, and such that a local maximum with respect to this neighborhood system is a global maximum for any particular right-hand side. In our example, the technology is simply given by the cost structure, and aside from non-negativity there is a single constraint requiring that output be greater than or equal to demand. We would have a larger number of constraints, even with this plant selection problem, if there were demands for output at several locations, or if we were explicit about a variety of factors of production. The neighborhood system would be independent of the demand specification and factor endowments. Figure 5 displays the neighborhood system for a problem with three, rather than two, types of plants.

In solving a series of programming problems in which the factor endowment changes from problem to problem, this minimal set of neighbors must be known explicitly or interrogated in some implicit fashion, since passing this quantity test is necessary for optimality. To the extent that a firm can be viewed partially as an algorithm for the solution of discrete programming problems, the firm must, in some fashion, be in possession of this set of vectors in order to test for optimality.

How, in general, are these neighbors to be determined for a given technology matrix? I find it astonishing that there is a canned computer program, that can be found either in *Mathematica* or *Maple*, which automatically calculates the set of neighbors if presented with the underlying activity analysis matrix. The program is not designed with this particular question in mind; its purpose is to compute a very sophisticated object in a field of mathematics known as Algebraic Geometry, a topic which is far removed from issues of economic theory. But here we see one of the remarkable, though rare, virtues of the translation into mathematical form of an every day problem: words, phrases, and concepts which bear no apparent relationship to each other in ordinary discourse may become synonymous in the language of mathematics.

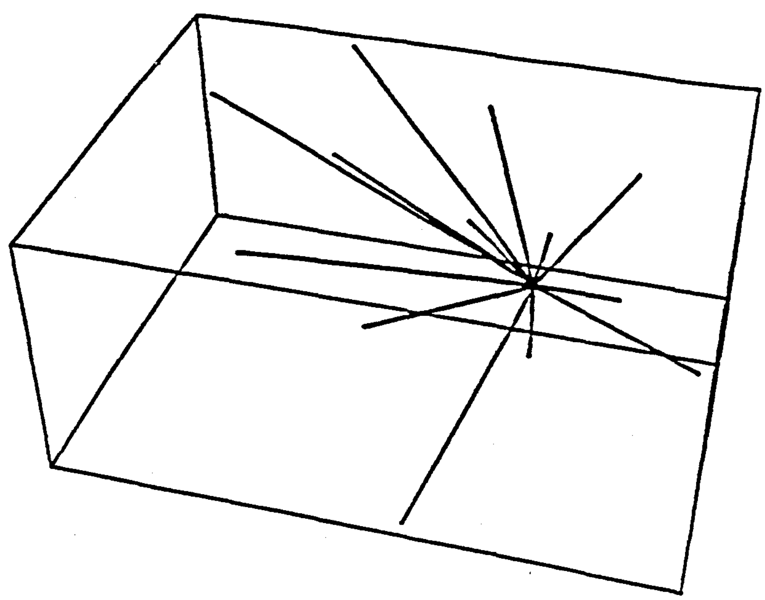

Figure 5

At the beginning of my talk, I spoke about comparative statics: about the analysis of changes in optimal behavior resulting from a modification in our economic environment. One type of modification is an exogenous change in factor endowments, or, in our example, a change in demand for output. The minimal neighborhood system permits us to analyze this type of change quite readily, in the sense that, for a general activity matrix, changes in the optimal solution associated with increases in the factor endowment or demand are given by precisely these neighbors.

A more complex change results from a modification in the technology rather than the factor endowment. Our numerical example is so elementary that the only changes in technology are essentially changes in the costs of the two competing types of plants. In order to see the consequences of such a change for the minimal neighborhood system, let us first make a proportional change in the parameters so that the capacities of the two plants are much larger.

	Smokestack	High Tech
Capacity	1600	700
Construction Cost	5300000	3000000
Marginal Cost	3000	2000
Average Cost	6312.50	6285.71

It is easy to see that the set of neighbors is unchanged by this rescaling even though the optimal solution for particular levels of demand will be quite different under the two regimes. Because of the increase in capacities, there will now be long intervals of demand in which the optimal plant configuration remains constant.

Now let us reduce the average cost at the Smokestack plant by raising its capacity to 1605, so that the parameters are given by:

	Smokestack	High Tech
Capacity	1605	700
Construction Cost	5300000	3000000
Marginal Cost	3000	2000
Average Cost	6302.18	6285.71

At this point, the average costs differ by only $16.47 and the minimal test set is increased by one new neighbor: it is now given by

$$(0,1),$$
$$(1,0),(1,-1),(1,-2),$$
$$(-1,3),(-2,5),(-3,7),$$
$$(4,-9),(7,-16).$$

As we see, the increased competitiveness of the two types of plants, associated with the convergence of average costs, requires a higher level of scrutiny in order to detect optimality.

For this class of problems, the set of neighbors is organized in a linear fashion, and small changes in the specification of the problem will always result in modifying our degree of resolution by adding or deleting an interval of neighbors at the end of the list. When the parameters change continuously, the unique procedure for detecting optimality changes in the most elementary fashion possible for a discrete, ordered set of points: the set grows or shrinks at one end. One of the major themes of my current research is to describe the ways in which the set of neighbors changes when the number of discrete choices is larger than two, and the neighbors are no longer organized linearly. All of the evidence that I have at the present moment suggests that, for the general integer programming problem, the minimal test set gains or loses members at a small number of locations on its boundary.

As a final example of a technical change, let the marginal cost at the High Tech plant rise to \$2015, so that the difference in average cost is only \$1.47. The set of neighbors becomes much larger:

Neighbor	Capacity	Cost
(0,1)	700	2310500
(1,0)	1605	5300000
(1,-1)	905	2989500
(1,-2)	205	679000
(-1,3)	495	1631500
(-2,5)	290	952500
(-3,7)	85	273500
(4,-9)	120	405500
(7,-16)	35	132000
(-10,23)	50	141500
(-17,39)	15	9500
(24,-55)	20	122500
(41,-94)	5	113000

It is certainly legitimate to suggest that this level of scrutiny is unnecessary; after all, hunting down a change of much less than 1/10 of one percent in average cost may be obsessive behavior bordering on pathology (even though total costs may be cut quite substantially by moving to a new plan). But let us leave this example with only two discrete choices concerning types of plants, and remember that in a large manufacturing enterprise there will be many discrete choices involving a large menu of tasks and machinery, each of which has its own capacity, set-up cost and marginal cost. The equipment may be placed in a number of different locations on the shop floor; the work may be passed from one piece of machinery to another with complex requirements of scheduling and precedence, and the tasks may alter from one job lot to another as the product specification varies. Demands may be revised capriciously and unexpectedly over time; output may be shipped to many different regions. The enterprise may have a host of competitors or none at all. In the absence of market prices regulating the flow of activity inside the enterprise, we cannot say in advance what the relationship between cost savings and detailed scrutiny may be and what degree of scrutiny will actually be required.

These examples also illustrate some unexpected structural elements of the set of neighbors: the set seems to be composed of a small number of linear segments. This is a very desirable feature, since the question of whether or not a member of a linear set of neighbors can be added to a proposed feasible solution so as to retain feasibility and decrease cost is easy; rounding will do. It is not difficult to argue that this structure is valid for an arbitrary problem with two integer variables; the set of neighbors always consists of a small number of intervals. This observation permits us to construct what computer scientists call a "polynomial" -- a really fast -- algorithm for this class of problems.

A remarkable accomplishment of mathematical programming is the generalization of this result to problems with an arbitrary number of integer variables. For any fixed number of integer variables, there is an integer programming algorithm which executes in "polynomial" time -- very rapidly -- as the other parameters of the problem vary. These algorithms have more than theoretical interest: they have been coded by experts, and seem to be among the best general purpose mixed integer programming algorithms currently available.

In summary, it seems to me that what I'm saying this afternoon boils down to a simple, straightforward piece of advice; if we are to study economies of scale and the division of labor in the large firm, the first step is to take our trusty derivatives, pack them up carefully in mothballs and put them away respectfully. They have served us well for many a year. But, derivatives are prices, and in the presence of indivisibilities in production, prices simply don't do the jobs that they were meant to do. They do not detect optimality; they aren't useful in comparative statics, and they tell us very little about the organized complexity of the large firm. Neighborhood systems are the discrete approximations to the marginal rates of substitution revealed by prices. They are relatively easy to compute, seem to behave pretty well under continuous changes in the technology and will ultimately lead to even better algorithms than we have now. We know much more about the structure of neighborhood systems than I have been able to describe this afternoon; not enough, perhaps, to derive a really satisfactory theory of the internal organization of the large firm at the present time. But my own intuition is that this is an important way to proceed; I am confident that serious, ultimately useful insights about the large firm will eventually be obtained by thinking very hard and long about indivisibilities in production.