Demand Theory and General Equilibrium: From Explanation to Introspection, a Journey down the Wrong Road

Alan Kirman[1]

Richard B Fisher Member,

Institute for Advanced Study, Princeton,

GREQAM, EHESS, Université d'Aix-Marseille lll

> Gerard Debreu symbolizes the use of a new mathematical apparatus, an apparatus comprehended by most economists only abstractly. Nevertheless, his work has given us an improved intuitive understanding of the underlying economic relevance. His clarity and analytical rigor, as well as the distinction drawn by him between an economic theory and its interpretation, have given his work important bearing on the choice of methods and analytical techniques within economic theory on a par with any other living economist.
>
> —*introduction of the Nobel laureate at the Royal Swedish Academy of Sciences,* 1983

The complexity of an economy stands in sharp contrast to the simplicity of a question that must be raised about its operation. Many agents compose the economy, and they have to deal with a large number of commodities. Each one of those agents makes decisions about the quantity of each one of those commodities that he will produce or consume: the number of variables involved is the product of the number of agents and the number of commodities. Moreover, in this decision-making process the agents act

[1] I would like to thank the Institute for Advanced Study for its hospitality and the participants of the HOPE conference for their helpful comments. In particular I have appreciated the valuable remarks of Phil Mirowski and John Chipman.

> independently of each other, and they are guided by self-interest. Why is high disorder not the result?
>
> The agents of an economy are counted in millions, if not billions. The number of commodities is similarly large. The self-interests of the independent decision-makers are sometimes in agreement, sometimes in conflict. Why does one not observe for every commodity a large excess of demand, evidenced for instance, by lengthy waiting times for orders to be filled, or a large excess of supply over demand, evidenced, for instance, by massive inventories?
>
> Agents no longer make independent decisions, and they interact with each other, if there are markets for commodities. Their interaction then reduces the difference between demand and supply.

—*Gerard Debreu, "Existence"* 1998

After the publication of the Arrow-Debreu model, economists were awestruck by the rigor and consistency of the reasoning. Finally, it was agreed, all the surplus flesh had been stripped off the skeleton of economic theory, and we now had new levels of analytic rigor to live up to if we were to be taken seriously as theorists. This was the evaluation of Gerard Debreu's Theory of Value (1959), in particular, and few, I think, were inclined to quarrel with the idea that the way forward was through modification and improvement of the general equilibrium model. After a long climb we had reached an intellectual plateau where, if one could breathe at such heights, everyone would like to be. Indeed, this idea fits well with a metaphor used by Einstein:

> Creating a new theory is not like destroying an old barn and erecting a skyscraper in its place. It is rather like climbing a mountain, gaining new and wider views, discovering new connections between our starting point

> and its rich environment. But the point from which we started still exists and can be seen, although it appears smaller and forms a tiny part of our broad view gained by the mastery of the obstacles on our way up."
>
> *quoted in Hahn* 1974

There were not many who would have argued that the mountain climbed by Arrow and Debreu was not the right one and that, in a relatively short time, even some of the stronger climbers would have reluctantly concluded that we were on a lonely peak far from the path traced out by earlier economists. What I will argue here is that the basic framework to which economic theory and, in particular, demand theory was reduced is inappropriate as a model of economic choice, since it depends on assumptions about individual agents derived not from observation but rather from introspection. Furthermore, these individuals are assumed to operate in an environment devoid of institutional content. Last, analytically the model contained the seeds of its own destruction. I would identify the crucial moment at which we committed ourselves to this rather futile path as coinciding with the appearance of The Theory of Value.

This route had three main features: first, the characteristics of economic agents; second, the environment within which they operate; and third, a notion of what would constitute an equilibrium state of a set of agents functioning in that environment. Suppose that, to simplify, we say that we are interested in the allocations of resources that are self-sustaining in some sense, and that such a state would constitute an equilibrium. By self-sustaining we might mean that nobody in the economy would, given the rules of the game, prefer to make another choice. That poses two immediate questions. Do such states exist, and will they be attained? But already I am jumping the gun. We must have a notion of what it means to prefer, and what the rules of the game are. Both of these have evolved along the path that I examine below and both help explain how we arrived at the wrong peak.

Before entering the heart of the matter, I want to indicate briefly what is involved here. The theory of individual preferences has been intensively developed, not because we are interested in the intrinsic qualities of those preferences but rather because we are concerned to model the choices that those preferences result in.[1] The various refinements and purification of preferences have been made with a view to maintaining certain properties of demand that, in turn, were necessary to guarantee the existence of equilibrium within the market framework I have alluded to. Thus preferences and demand have been intimately linked. The second remark relevant to the whole of what follows is that the market as an institution hardly figures in the theoretical literature. Indeed, as Douglas North (1977, 710) says, "It is a peculiar fact that the literature on economics . . . contains so little discussion of the central institution that underlies neoclassical economics—the market."

For the moment, let me briefly mention what is meant by market in the history of general equilibrium theory. I will then discuss the difficulties that have arisen with this choice of framework and then come back to the importance of taking an alternative view and an alternative path.

When Debreu talks about markets, all that he has in mind is that the prices of all goods at all times and in all places should be determined and known to all participants in the economy. This vector of prices constitutes what he refers to as "market prices." In particular, he is interested in those prices for which the excess demand is zero for every commodity, and it is these that he refers to as "equilibrium prices," and in this he follows exactly in Léon Walras's footsteps. On the subject of how and by whom those prices are established, he is silent, as he is on the subject of how the trades that correspond to those prices are effected. Yet there is the illusion that such considerations are present. To be more specific, consider the last part of the second epigraph, which reveals why the model proposed does not answer the question posed. Debreu says, "Agents no longer make independent decisions, and they interact with each other, if there are markets for commodities. Their interaction then reduces the difference between demand and supply." But this is precisely the basic weakness of the model. Nowhere is there any explanation of how this happens or, beyond what I have already said, what is meant by "if there are markets for commodities." Yet these problems were central to the questions posed by Walras, for example, as Donald Walker (1996) has pointed out. Indeed, at first glance, it seems that Debreu is alluding not only to the equilibrium but also to an adjustment mechanism when the prices are not at their equilibrium values. Yet all who are familiar with Debreu's work know that, unlike Kenneth Arrow, Frank Hahn, and Leonid Hurwicz, Debreu was not preoccupied with the problem of how prices might adjust to equilibrium. He was concerned with the existence of equilibrium and only in a very tangential way, through his work on local uniqueness, with its stability.

By the time that we have arrived at the peak first climbed by Arrow and Debreu, the central question boils down to something rather simple. We can phrase the question in the context of an exchange economy, but producers can be, and are, incorporated in the model. There is a rather arid economic environment referred to as a purely competitive market in which individuals receive signals as to the prices of all goods. All the individuals have preferences over all bundles of goods. They also have endowments or incomes defined by the prices of the goods, and this determines what is feasible for them, and the set of feasible bundles constitutes their budget set. Choosing the best commodity bundle within their budget set determines their demand at each price vector. Under what assumptions on the preferences will there be at least one price vector that clears all markets, that is, an equilibrium? Put alternatively, can we find a price vector for which the excess demand for each good is zero? The question as to whether a mechanism exists to drive prices to the equilibrium has become secondary, and Herb Scarf's famous example (1960) had already dealt that discussion a blow.

The warning bell was sounded by such authors as Donald Saari and Carl Simon (1978), whose work gave an indication, but one that has been somewhat overlooked, as to why the stability problem was basically unsolvable in the context of the general equilibrium model. The most destructive results were, of course, already there, those of Hugo Sonnenschein (1974), Rolf Mantel (1974), and Debreu (1974) himself. But those results show the model's weakness, not where that weakness comes from. Nevertheless, the damage was done. What is particularly interesting about that episode is that it was scholars of the highest reputation in mathematical economics who brought the edifice down. This was not a revolt of the lower classes of economists complaining about the irrelevance of formalism in economics; this was a palace revolution.

How did we wind up on this peak, and why has it been so difficult to leave it? On my bookshelf sit several volumes of Bourbaki's monumental Elements de mathematiques. Not far from them is a book edited by, among others, John Chipman. And nearby is Debreu's Theory of Value. The three epitomize a view of theoretical economics that flourished from the sixties to the seventies. Economics was, at least in the view of those in mathematical economics at the time, ready to build a structure as minimalist and abstract as that provided by Bourbaki for mathematics. It is the emergence of this phenomenon that I would like to explain. The first element is the introduction of the axiomatic method, something in which Debreu was a fervent believer and which was certainly not unrelated to his own mathematical education in the shadow of the Bourbaki crowd. This did not simply reduce the existing framework but changed its orientation (see Mirowski and Weintraub 1994). The second is the nature of the preferences, which are held to characterize economic agents. If one looks back to that classic book edited by Chipman et al. (1971), one sees how far we have moved in our vision of what constitutes rational choice. It would be difficult to imagine a more erudite, arid, and inaccessible volume than theirs, yet it was widely considered as exemplary scholarship. There was little discussion of the underlying foundations of the very notion of preferences but rather a series of contributions trying to establish various relations between preferences and demand. On what preferences are, Hirofumi Uzawa (1971,) says simply,

> Preference relations, in terms of which the rationality of human behaviour is postulated, are precisely defined as irreflexive, transitive, monotone, convex and continuous relations over the set of all conceivable commodity bundles. A demand function associates with prices and incomes those commodity bundles that the consumer chooses subject to budgetary restraints.

Uzawa's basic and important contribution was to show that if a demand function satisfies the weak axiom of revealed preference (WARP), then there exists a preference relation from which it is derived. Furthermore, demand functions derived from preferences satisfy WARP. Both Jean Ville and Hendrik Houthakker (see Gardes and Garrouste, this volume) had shown that a necessary and sufficient condition for a demand function to be derived from preferences was that the strong axiom of revealed preference (SARP) be satisfied. Uzawa completed the picture by showing that, under certain regularity conditions, WARP implies SARP. This is more than a little disappointing, since it was originally thought that revealed preference theory was somehow more satisfying than utility-based theory, since it is based, at least in principle, on observable behavior. Thus, with these results, an avenue out of the abstract world of preference relations satisfying certain axioms derived from introspection was closed off.

One might have thought that revealed preference would have lost its appeal after this result, but this was not the case. The reason for this is clear. Although demand choices are in fact observable (in principle) in the Arrow-Debreu framework, an individual is never faced twice with the same situation. Choices are made once and for all at the beginning of time. Since all commodities are dated, seeing an individual choosing bundle x when she could have chosen y; then, subsequently choosing y when she could have chosen z; and then later, choosing z when she could have chosen x is not a violation of WARP, since the x that figures the second time is not the same as the one in the first choice. The commodities are available at a later time and are therefore no longer the same. Furthermore, if any consumption took place between the choices, the preferences that determine the later choices may well be affected. Therefore, if we adhere to the general equilibrium framework, we can never actually observe a contradiction with the hypotheses that characterize revealed preference. The only way out of this would be to pose the hypothetical question repeatedly to an individual: which of these two bundles would you choose? At this point the advantage of the "observability" of the choices has been lost. Yet revealed preference still plays a prominent role in theory and in the teaching of economics (see Mas-Colell, Whinston, and Green 1995). It is prompted by the simplicity with which it can be explained, and this is reinforced by its widespread use in experimental economics. In the experimental laboratory the practice of posing successive hypothetical choice questions and then checking to see whether these choices are coherent is widespread. There are technical questions, such as whether to prefer Sidney Afriat's (1967) strong axiom, which is convenient when one is faced with a finite set of observations. But the underlying logic remains: what we are demanding of our unsuspecting experimental subjects is no more or less than what we require of our traditional economic agents endowed with preferences with all the standard properties.

Moreover, a great deal of effort and mathematical skill were spent on weakening the conditions imposed on preference relations. It was argued that weakening the assumptions in this way would make them more acceptable. Little was said about the basic underlying assumption that individuals are indeed characterized by such preference orders. Taking this as given, the goal was to weaken the basic axioms as much as possible. Today, it is recognized that none of the standard axioms, even in their weakened form, are derived from observation of choice behavior but rather are the result of pure introspection by economists. Yet this recognition has nothing to do with recent developments. Werner Hildenbrand (1994) cites quotations from economists such as Lionel Robbins (1935), Tjalling Koopmans (1957), and John Hicks (1956), all of whom were well aware of this. Worse, later in his life, Vilfredo Pareto concluded that individuals make quite arbitrary choices and spend the rest of their time justifying them! Thus skepticism about the bases of the utility theory on which the theory of demand is based is long-standing among highly reputable economists.

If one does embark on this theoretical route, it seems natural to make the assumptions as weak as possible. But Pareto himself was not happy with this insistence on weakening assumptions to the maximum. He was still, for example, convinced that the idea of measurable utility had some merit. As Chipman (1973) pointed out, "Thus altering slightly a suggestive metaphor of Georgescu Roegen, just because the equilibrium of a table is determined by three of its legs, we are not required by any scientific principle to assume that actual tables have only three legs, especially if direct observation suggests that they may have four." Other economists on the same road were less reticent. Hicks, for example, suggested that the principle of Occam's razor should always apply, and he was far from alone. Indeed, as we move down the road from Walras and Pareto we observe a clear tendency to do this, and the driving force is not economic plausibility but mathematical convenience. With the formalization and axiomatization of the basic model, it became clear that what was required for equilibrium to exist was some sort of continuity of demand and hence of excess demand. This was what pushed economists toward assumptions of convexity and continuity. If the continuity of demand could be guaranteed by weaker assumptions than the standard ones, this was welcomed. Thus the name of the game was to find the weakest conditions under which equilibrium exists.

Why this insistence on the existence problem? Simply because given how the program was formulated, the natural idea was to look at "states of the economy" and then to characterize the equilibrium states. This preoccupation was not new, but it came sharply into focus in the Arrow-Debreu setting. When economists talk about the state of the economy, they may mean different things. Depending on the particular problem at hand, the economist may be interested in a detailed description of some part of the economy or may simply wish to examine some highly aggregate variables. At the most general level we might be interested in specifying everything about the allocation and activity of every agent in the economy, and, indeed, this was the lofty ambition of general equilibrium theory. At the other end of the spectrum we might simply be interested in the number of people "out of work" or the "price level." Theoretical micro economists have typically been interested in the first problem and have been preoccupied since the early 1900s with, for example, the characteristics of the competitive equilibrium of a fully specified economy.

It is important to recall at this point that all of this theory is developed in the context of the perfectly competitive system. In the general equilibrium model, no "organization" or specification of interpersonal relationships is assumed. The consistency of allocations in the sense that what individuals receive is just what they desire is achieved by an anonymous price system that all agents take to be given. Thus there is no need for any specification of the relationships within the economy; the only way in which agents communicate is through the price system. Indeed, to pursue their analyses, economists have stripped away all the institutional details and structures that were still very much present in the work of both Walras and Pareto and even more so in the work of Alfred Marshall. In particular, economists have focused on the assumptions about individual characteristics and resultant behavior, which will ensure achieving an "equilibrium outcome."

This is, of course, a very static view of economics. Even when looking at economic dynamics, economists have, in general, concentrated on looking at "steady states" or equilibrium growth paths. How did economists become so wedded to this view of the world? A neat and clear explanation of the development of economic theory in this century would be that it was essentially dominated by the analytic methods of classical mechanics and that over most of the last century, physics triumphed as a model for economics. This is indeed the view advanced by Bruna Ingrao and Giorgio Israel (1990) and Philip Mirowski (1989). In other words, the peak, which we reached in the Arrow-Debreu model, was indicated by signposts from physics, which have been there since the nineteenth century. Yet, it seems reasonable to ask, is this explanation justified?

Economics and Physics

What is clear from both Pareto's analysis and that of many of his contemporaries such as F. Y. Edgeworth, William Stanley Jevons, and Franklin Fisher is that they all shared a conviction that there was an analogy between economic systems and those of classical mechanics. Edgeworth was quite explicit in suggesting that a "mecanique sociale" would take its place alongside the "mecanique celeste" (quoted in Cohen 1994, 90). Jevons (1905, 631) said that economics resembles physics in that "the equations employed do not differ in general character from those which are really treated in many branches of physical science." Another contemporary, John Elliott Cairnes was even more explicit: "Political Economy is as well entitled to be considered a 'positive science' as any of those physical sciences to which this name is commonly applied" (quoted in Cohen 1994, 42). He went on to argue that the principles of economics have identical features to those "of the physical principles, which are deduced from the laws of gravitation and motion" (198). Mirowski (1989), Ingrao and Israel (1990), and Bernard Cohen (1994) have examined at length the validity and consequences of such assertions.

The extent to which the analogy between physics and economics has ensnared economics in a position that it could have avoided had it found its source of inspiration elsewhere—for example, in biology, as Marshall suggested—is well documented by these authors. Pareto (1953, 185) himself remarked that, when examining the equations that have to be solved to determine an economic equilibrium, someone well versed in mathematics or physics would say, "These equations do not seem new to me, they are old friends. They are the equations of rational mechanics." He went so far, in the Cours, as to draw up a table of analogies between the two disciplines.

Although this argument is persuasive and allows us to establish a simple route from Walras to Debreu, it is not completely convincing. The vision of equilibrium is subtly different in the two cases. In economics it is really the solution of a system of equations, a fixed point; in the physics interpretation, at least the metaphor of a resting point of a dynamical system persists as, indeed, it did in the writings of Walras and Pareto.

In the economics view, agents in an economy maximize some concave or quasi-concave objective function over a convex set defined by parameters over which they have no control. The natural way to analyze such a system is to look for an equilibrium state, that is, values of the parameters determining the constraints of the individuals and choices of those individuals such that none of them has any incentive to modify his or her actions. The passage from individual to aggregate behavior is simply additive. Furthermore, no explanation is given as to how the equilibrium comes about; this is the subject of stability analysis, which has usually been reduced to examining the convergence of artificial price-adjustment processes. Yet, as Michio Morishima (1964,) pointed out,

> If economists successfully devise a correct general equilibrium model, even if it can be proved to possess an equilibrium solution, should it lack the institutional backing to realise an equilibrium solution, then the

equilibrium solution will amount to no more than a utopian state of affairs which bear no relation whatsoever to the real economy.

Uniqueness and Stability of Equilibrium

As I have said, there is literally nothing in the Arrow-Debreu model about how an economy out of equilibrium could attain that state. If we are interested in that problem, then we must show how a nonequilibrium price vector would be altered toward or into equilibrium. The simplest adjustment process is the well-known tâtonnement process suggested by Walras, which he argued would lead naturally to equilibrium. He assumed that the level of prices for which there is a positive excess demand should rise and the level for those for which the excess demand is negative, that is, for which there is excess supply, should fall. The tâtonnement process can be written as follows:

$$p_{t+1} = p_t + \lambda Z(p_t)$$

Just writing this expression presents a problem. What is the time involved? The usual argument is that one should think of the adjustment as taking place in virtual time. It is common practice also to write this process in continuous time as

$$\dot{p} = \lambda Z(p)$$

Since the time involved is not related to the model, this does not really matter. Yet this is purely artificial. Although it is easy to convince oneself of the basic intuition behind the process, the way in which it works is very odd. Indeed, it usually assumed that no trade takes place until equilibrium prices are established.

Whenever one looks closely at the mechanics of the adjustment mechanisms that have been proposed, however, they almost all refer back to Walras. The economists just mentioned saw themselves as being the direct descendants of Walras and Pareto. Hicks

(1939, 2), for example, says, when he wishes to spell out the adjustments on several markets,

> When looking for such a technique, we are naturally impelled to turn to the works of those writers who have specially studied such interrelations—that is to say the economists of the Lausanne school, Walras and Pareto, to whom, I think, Wicksell should be added. The method of General Equilibrium, which these writers elaborated, was specially designed to exhibit the economic system as a whole, in the form of a complex pattern of interrelations of markets. Our own work is bound to be in that tradition, and to be a continuation of theirs.

Despite the contribution of Edgeworth, the "written pledges" model of Walras, and later brave efforts by Hicks, Hahn, Takashi Negishi, and others, analysis of so-called non-tâtonnement processes, in which more realistically economic activity is allowed before equilibrium is reached, has received much less attention. The desire to pursue a more realistic mechanism through which market equilibrium could be attained led Hicks to establish the well-known "Hicksian week" during which trade would take place at nonequilibrium prices. In a sense, this is an alternative to Walras's "written pledges" model, described in detail by Walker (2005). Yet in all the non-tâtonnement adjustment processes proposed, there is always some artificial argument about the consistency between the sign of individual and aggregate excess demands as the process proceeds. In effect, "nobody is on the wrong side of the market." This is disappointing, since non-tâtonnement is clearly more realistic than tâtonnement. However, let us go back to the route that was followed and leave non-tâtonnement processes to one side, as the literature did, and assume for a moment that we accept the artificial tâtonnement process.

If such a process did actually converge to equilibrium from an arbitrary starting point, could it be said that things would surely be much simpler? Not really, because the fundamental problem is that the conditions, which are known, which guarantee the stability of this process, cannot be obtained from assumptions on the behavior of the individuals. To be absolutely clear, what Sonnenschein (1972), Mantel (1974), and Debreu (1974) showed is that there is no hope of a general result for stability, since the only conditions on the aggregate excess demand function that can be derived from even the strongest form of the assumptions on individual preferences are the well-known four: continuity, Walras's law, homogeneity of degree zero, and boundary conditions that guarantee that aggregate excess demand "explodes" if any price goes to zero. Since there are functions satisfying these conditions that are not stable with respect to the tâtonnement process, and that do not have a unique equilibrium, these properties cannot be guaranteed by the assumptions on individual characteristics.

The full force of the Sonnenschein, Mantel, and Debreu (SMD) result is often not appreciated. Without stability or uniqueness, the intrinsic interest of economic analysis based on the general equilibrium model is extremely limited. Morishima's observation about stability is well taken, but for macroeconomists uniqueness is also important. "Comparative statics" in which one compares equilibrium and another one, predicated on a change in the parameters, makes no sense in the presence of multiple equilibria. The usual way out of this problem is to assume a "representative agent," and this obviously generates a unique equilibrium. However, the assumption of such an individual is open to familiar criticisms (Kirman 1992; Stoker 1995), and recourse to this creature raises one of the basic problems encountered on the route to the place where general equilibrium has found itself: the problem of aggregation. In fact, we know that, in general, there is no simple relation between individual and aggregate behavior, and to assume that behavior at one level can be assimilated to that at the other is simply erroneous.

At this point it is worth mentioning an attempt to restore some structure that had been abstracted away and, hence, to overcome the difficulties produced by the fact that uniqueness and stability of equilibria cannot be guaranteed in the basic model. Interestingly, this stems from ideas advanced already by Augustin Cournot. This approach suggests that, if the economy consists of a large number (necessary for price-taking behavior to make sense) of sufficiently heterogeneous agents, then properties like uniqueness and stability of equilibrium may be restored (Grandmont 1987, 1992; Hildenbrand 1983, 1994). Structure may be introduced into aggregate behavior by the presence of sufficient differences between the characteristics of the agents. There is, of course, an important deviation from the conventional route here, the idea being to introduce assumptions on the distribution of preferences rather than on the individual preferences themselves. Unfortunately, this approach has not yet led far, as recent work by Etienne Billette de Villemeur (1998) and Hildenbrand and Alois Kneip (2005) has shown. The problem seems to reside with the definition of heterogeneity. What has been shown so far by Hildenbrand is that heterogeneity of behavior because of particular income distributions imposes structure, a version of the "law of demand," on aggregate behavior. What has not been shown, as yet, is how one can make assumptions on the distribution of the underlying preferences in the standard model that will guarantee the dispersion of behavior. It is conceivable that dispersion of the parameters of utility functions may well not lead to different behavior of the individuals. The increasing dispersion introduced by Grandmont essentially makes the economy behave as if it were one giant Cobb-Douglas individual and the representative agent appears again. However, my basic argument here would be that the difficulties with the basic model are too fundamental to be solved by such assumptions on the distribution of characteristics.

Another argument that has been used to diminish the impact of the SMD result is that of Donald Brown and Rosa Matzkin (1996). They argue that excess demand should be thought of as depending on both prices and income, and that doing this will impose some structure on aggregate excess demand. To take the simplest example: suppose that we bound endowments and then insist that consumption shall not be negative; then we may have restrictions on possible excess demand functions. At this point the argument becomes almost philosophical. We know from an early result of Alan Kirman and Karl Koch (1986) that we can posit a function and an income distribution and still find an economy that corresponds to it. However, what is also true is that, to do this, we cannot bound the resources a priori. Chipman (this volume) gives an elegant account of related problems. Few will, I think, find that the basic message of the SMD result has been really weakened. As Hahn (2002, 224) has written, when talking of the attempts to establish stability results and referring to the SMD result, "The enterprise was doomed not to be capable of reaching general conclusions in the Walrasian setting. A theorem not directly related to connected with dynamics did the damage." Having examined the weakness of the general equilibrium model and its incapacity to tell us anything about the functioning of the economy out of equilibrium, we might ask whether we cannot establish the source of the difficulty.

Information and Stability

One of the major attributes of general equilibrium theory also reveals its Achilles' heel. It is often asserted that one of the major achievements of general equilibrium theory is its ability to demonstrate that the competitive mechanism is remarkably parsimonious in terms of the amount of information that it uses. Such an observation is couched in terms of the mechanism design literature because of Hurwicz and others (for a discussion of the mechanism literature, see Lee, this volume). A mechanism for attaining allocations for a given economic environment specifies the messages that have to be transmitted by the economy and then a mapping that translates those messages into allocations. The measure of how much information such a mechanism uses is considered to be the dimension of the message space. Suppose that we are interested in attaining Pareto efficient outcomes. We know that the outcomes of the competitive process achieve this. The standard theorem—derived, in its most general form, from Jordan 1982—says that the competitive mechanism needs a "message space" with dimension $n(l-1)$, where n is the number of agents and l the number of goods. This is quite remarkable. To see why this is so, consider a simple exchange economy. At equilibrium every agent except one (since at equilibrium it is enough to know the excess demands of n-1 agents) needs to transmit his vector of excess demands ($l-1$ goods because of Walras's law), and, thanks to the homogeneity of degree zero of excess demand functions, a vector of $l-1$ prices is needed. This result seems, at least at first sight, to justify the claims for the efficiency of the competitive market mechanism, since one can show that no other mechanism that achieves Pareto efficiency uses less information. This result can be extended to include cases in which we are interested in attaining allocations that satisfy other criteria.

For example, consider a very old problem in economics, that in which one is interested in mechanisms that not only will achieve Pareto efficiency but will also be "fair" in the sense that no individual prefers someone else's allocation to his or her own. This would seem to be a very demanding requirement in terms of information, since every individual must be able to compare his or her allocation with that of all the other agents. Nevertheless, one can show that a mechanism with a message space of only nl dimensions can achieve this result (Calsamiglia and Kirman 1993). In other words, the information required is hardly more than that used by the competitive mechanism to achieve Pareto efficient outcomes. The secret of this result is to use the competitive mechanism from an equal division of resources. Thus the Walrasian mechanism permits an enormous economy of information. One might argue that this is a strong reason for thinking that we have not come down the wrong road, and that those who defend the modern version of general equilibrium analysis are justified in doing so. In my view this is wrong; to see why, one has to recall that the results just mentioned show only that little information is needed for the economy to function at equilibrium. But what this says is that if the economy is actually at an equilibrium, the amount of information needed for it to function is very limited indeed.

But what should interest us as economists is not only how informationally demanding the mechanism is at equilibrium but also how much information it requires to get there. We find ourselves back at the problem of stability, one that, and it is worth insisting on this again, Debreu consistently avoided and indeed later claimed that he had always considered to be beyond hope. Yet, as I have already said, it is surely the case that economic equilibria are of interest only if they can be attained through a reasonable adjustment process. So if we want to ask the question as to how much information is involved, we have to rethink the original problem in terms of finding an adjustment process that will modify the messages and be guaranteed to achieve Pareto efficient outcomes. As we have seen, the tâtonnement process from an economic point of view has a basic problem. The SMD result shows that the equilibria of economies are not necessarily stable with respect to that process. It is then natural to argue that the problem lies with the adjustment process rather than with the general equilibrium model. If a more general adjustment rule were to be specified, perhaps the equilibria of the economy could be shown to be stable. Yet what became immediately clear after the innovative work of Steve Smale (1976) was that stability could be achieved only at the price of an enormous increase in the amount of information required.[2]

Smale's global Newton method extends standard methods to allow one to find a fixed point of a mapping, such as an aggregate excess demand function, if one starts sufficiently near the boundary of definition. It has two major drawbacks. First, it does not behave well in the interior of the domain that, in the case under consideration, is the space of all strictly positive prices. Second, as already mentioned, it uses a great deal of information. What is needed, at each point in time, is knowledge of all the partial derivatives, and this greatly increases the size of the message space, without guaranteeing convergence from any arbitrary starting point. An additional objection would contend with the economic content of the process. While the original tâtonnement process has a natural interpretation, this is not the case, despite the efforts of Hal Varian (1977), for the Newton methods.

Do the informationally demanding requirements of the Newton method represent a necessary evil? Saari and Simon (1978) asked the following question: can one find what they called “locally effective price mechanisms,” that is, ones that turn all economic equilibria into sinks or attractors and that use less information than the Newton methods? They proved, unfortunately, that this cannot be done. This is odd, since the generalized Newton method has the undesirable and restrictive property that it reduces excess demands monotonically and that any single market in equilibrium is kept in equilibrium. One might have hoped that, by relaxing this requirement, one could have found less informationally demanding mechanisms.

As Jim Jordan (1986) pointed out, all the alternative adjustment processes that had been constructed at the time he wrote had no natural economic interpretation. There have been many efforts to construct globally and universally stable price-adjustment processes since then, and in a certain sense Kazuya Kamiya (1990), Gerrit van der Laan and Dolf Talman (1987), Peter Flaschel (1991), Jean-Jacques Herings (1995), and Jan Tuinstra (2002) have succeeded. Yet if one looks closely at these results, there is always some feature open to objection. In Kamiya’s case the excess demand function is artificially defined outside the original price domain. In Flaschel’s case the adjustment process depends on a parameter that varies with the economy, and, indeed, he says that it is too much to hope that one would find a process that would work for all economies. Herings’s mechanism has the curious feature that prices are adjusted according to the relation between current price and the starting price. All of this seems, to me at least, to suggest that there is no hope of finding an economically interpretable