Introduction

The papers in this volume extend from my first published paper (Chapter 12) to very recent work. Generally, they are papers in theoretical and applied microeconomics. As in previous volumes (Fisher, 1991, 1992a, 1992b), I have not reprinted papers that have been superseded by my other books.

The chapters of Part I ("Disequilibrium and Stability"), particularly those of Part IA ("Models of Disequilibrium Behavior"), reflect my principal theoretical interest of the 1970s and early 1980s other than as discussed in my book on the subject (Fisher, 1983). In approximately 1970, I became convinced that the single most important lacuna in economic theory is our inability to explain whether or how competitive economies not in general equilibrium succeed in getting there. In part under the (usually) benign influence of Frank Hahn, who visited MIT in 1971, I explored that interest for well over a decade.

I remain convinced that the problem of disequilibrium and stability is one of central importance. Briefly, the elegant nature of the economics of equilibrium, beginning with the analysis of the plans of optimizing agents, has caused economists to concentrate on situations in which those plans are mutually compatible. But the work of the Invisible Hand cannot be understood by looking only at situations in which that work has been completed. Further, the central tenets of Western capitalism as to the desirability of free markets – the relations between competitive equilibrium and Pareto efficiency – become empty propositions if positions of competitive equilibrium cannot be achieved or can be achieved only slowly. These propositions are further discussed in Chapter 1, which also summarizes both the literature and my work as of the mid-1980s. It will be seen that the state of the art remains unsatisfactory. This may be both a consequence and a cause of the phenomenon that the profession largely continues to ignore such issues or, at least, to behave as though they had long been satisfactorily resolved.

This is not only a matter of high-powered general equilibrium theory.[1] We do not even have a really satisfactory analysis of how competitive firms set prices in a single market, since each firm is supposed to take prices as given. I therefore began my study of these issues by examining the problem of price change in an oversimplified setting. That resulted in Chapter 2, which deserves its subtitle of "A Preliminary Paper." In that paper, I tried to formulate a model in which more or less competitive firms set prices by hunting for the "correct" competitive price. To do this, I had to require that not all information was perfect and instantaneous, and the obvious choice was to allow customers to search among firms in an attempt to find the firm with the lowest price. Each firm was then assumed to adjust prices according to whether it sold more or less than its planned supply. It was not very hard to find assumptions under which this process converged to (partial) competitive equilibrium.

Chapter 2 was written in the early days of what was to become the extensive literature on search processes. Indeed, it was written at about the same time as the well-known article by my colleague, Peter Diamond (1971), which, in a somewhat similar model, showed how a search process could end up at a monopoly price. The difference in conclusion stemmed from a difference in purpose. Diamond quite reasonably set up a search model and found out where it led. I, on the other hand, was beginning a research program to discover whether relatively plausible disequilibrium stories could have a competitive ending. If not, then the underpinnings of microeconomics would be shaky indeed. I therefore built a model designed to have such an ending.

Unfortunately, that construction could not be regarded as fully satisfactory. As Michael Rothschild (1973) pointed out in his review of the early search literature, the firms described in Chapter 2 do not behave very sensibly. They take their demands to be independent of price even though they can readily observe that their sales are higher at lower prices than at higher ones. Chapter 3 was an attempt to remedy this failing by building a model in which the market power that disequilibrium conferred on firms asymptotically disappeared. This proved remarkably difficult to do in any very convincing way, and the question of how (or even if) one gets to equilibrium in a partial setting remains without a rigorous theoretical answer.

The fact that disequilibrium leads to perceived market power and that

[1] Indeed, it is not a matter of theory only. The habits of equilibrium analysis and the failure to think about adjustment processes infect the way in which economists analyze real-life phenomena and the policy recommendations that they give. This extends from macroeconomics and rational expectations to the analysis of antitrust cases (see Fisher et al., 1983).

this influences the path and possibly the equilibrium of the system was to reemerge, this time in a general disequilibrium context. As discussed in Chapter 1 in my book on disequilibrium (Fisher, 1983), I showed that one could model price offers based on perceived market power. Further, the question of whether the economy reaches a Walrasian equilibrium or remains transaction-constrained depends directly on what happens to such perceptions. Indeed, the question of whether agents are constrained by the need for money also depends on such matters. In a sense, the two revolutions in economic theory of the 1930s – the Keynesian possibility of the liquidity trap and non-Walrasian equilibrium on the one hand and the introduction of imperfect competition on the other – turn out to be related.

Before Chapter 3 was written, I had turned my attention directly to the problems of disequilibrium and stability in a general competitive setting. The first fruit here was Chapter 4.

As described in Chapter 1, stability theory began with tâtonnement, where a fictitious auctioneer adjusts prices in the direction of excess demand while nothing else happens. After a promising beginning, however, the hope that this (in any case hopelessly unrealistic) model would generally lead to stability was dashed in 1960 by the discovery of a class of counterexamples (Scarf, 1960). At about the same time, new models were being developed. These new "non-tâtonnement" or "trading processes" did not do away with the auctioneer but permitted trade as well as price change to take place in disequilibrium. (Despite the powerful results that were obtained, for many years much of the profession continued to believe that stability theory concerned only tâtonnement.)

Two basic models of trading processes were developed. One of these was the Edgeworth Process in which trades were assumed to take place if and only if a group of agents could improve themselves by trading at constant prices.[2] The second was the Hahn Process, whose driving assumption was that markets are organized so well that, after trade, unsatisfied buyers and unsatisfied sellers of the same commodity cannot both remain.

For reasons discussed in Chapter 1, I regarded (and still regard) the Hahn process as the one likely to lead to really satisfactory models, and, in the early 1970s, I began to write a series of papers designed to explore and extend it. In large part, this led to my 1983 book, in which the analyses of the relevant papers are given and improved; however, Chapters 4 and 5 were not entirely subsumed in that volume.

[2] For a more detailed discussion and bibliography see Chapter 1 of this volume or Fisher (1983).

Chapter 4 is the more important of these two pieces. It represents an attempt to do away with the fictitious auctioneer.[3] I observed that if one regards goods sold by different sellers ("dealers") as different goods, then the Hahn Process assumption becomes truly compelling. Of course, such a treatment permits dealers to make price offers, and here I used the unsatisfactory device of Chapter 2 in which each dealer simply adjusts his or her price according to whether sales are greater or less than planned. One must also do something to ensure that, in equilibrium, the prices charged by dealers in the same good come together. This approach still seems to me to be a useful one, but it was only very partially continued in my later work.

As noted, that work mainly resulted in my earlier book (Fisher, 1983). There I tried to analyze what is really the main problem: Does an economy in which agents recognize and act on the arbitrage opportunities created by disequilibrium tend to converge? The answer turns out to be difficult, and one might reasonably say that the major interest of the book lay in the study of disequilibrium itself rather than in the stability theorem proved. Chapter 1 summarizes the results reached.

Chapters 5 and 6 represent two alternate attempts to model the results of disequilibrium awareness. In Chapter 5 (written before the work leading directly to my earlier book), agents are assumed to discover constraints on their trading and to optimize taking those constraints into account.[4] Queues form, and some agents do not get served. It is shown that so long as (after trade) agents with an unsatisfied excess demand have that demand of the same sign as their unconstrained ("target") demand, then certain stability properties follow. Indeed, under fairly plausible assumptions, the only rest points of the model will be Walrasian, even though one begins with constraints on trade. The assumptions, however, make sense only in the neighborhood of Walrasian equilibrium.

The model of Chapter 5, like the work in my 1983 book, is based on the Hahn Process. Chapters 6 and 7, by contrast, deal with the Edgeworth Process. I cannot say that my 1983 book aroused very widespread interest – the profession continuing its practice of overlooking such issues – but some people were very interested indeed (among them, Maarten-Pieter Schinkel, the editor of the present volume). One of these was Dale Stahl II, who visited MIT in the mid-1980s. He and I combined our work

[3] Incidentally, the remark in Chapter 3 that this is like trying "to play *Hamlet* without the Prince of Norway" is not a slip. The role of Fortinbras in *Hamlet* is that of one who is extraneous to the action but cleans everything up.

[4] In neither Chapter 5 nor Chapter 6, however, do agents understand that prices will change.

and wrote Chapter 6. That paper presents an alternate model in which agents are aware of disequilibrium. In it, unsatisfied demand is rationed by a system of queues. The result is a stability theorem based on the Edgeworth Process.

Despite this foray, I continue to believe that the Edgeworth Process is not an attractive model for out-of-equilibrium trading. Indeed, the simple assumption on which it is based is not nearly as appealing as it first appears. One of the reasons for this is explored in Chapter 7. (Other reasons are given in Chapter 1.) Chapter 7 has a mildly interesting history. In 1975, I was invited to give the F. W. Paish Lecture at the meeting of the Association of University Teachers of Economics in Sheffield, England. I thought it was time to reflect on the state of the stability literature and presented an address that was later published as Fisher (1976) and (somewhat revised) as Chapter 2 of my (1983) book on disequilibrium. In doing so, I commented that one of the problems with the Edgeworth Process assumption is that it might require very large numbers of agents to find each other in order to produce a mutually improving trade at given prices. Reflecting on this, Paul Madden produced a paper (1978) in which he showed that, provided every agent always held a positive amount of every commodity, the trades involved could always be bilateral. At about the same time, David Schmeidler (privately) pointed out to me that the number of agents required for an Edgeworth Process trade need never exceed the number of commodities.

I felt that neither of these results, while interesting, really provided a satisfactory answer to my objection. Particularly if one dates commodities, the number of commodities is likely to exceed by far the number of agents, so Schmeidler's bound does not seem very helpful. Moreover, the assumption that every agent always holds a positive amount of every good is far too strong to be sensible, particularly in a model of disequilibrium. But it was not until the late 1980s that I tried to explore the consequences of relaxing that assumption.

The result was the present Chapter 7, in which I show that the number of agents that might have to be involved in an Edgeworth Process trade is substantial. Further, some of the required trades could be quite complicated indeed, involving circumstances in which certain agents were induced to make trades that they did not want in order to induce others to make trades that the original agents found desirable. Some years later, one of my undergraduate students, A. D. Tsai, proved (in an as-yet unpublished paper) that the problem of finding such trades is NP-hard, so that the assumption that trade takes place whenever a group of agents can improve themselves by it is by no means a weak one.

It is evident that what is involved here is a basic question about the way in which markets come into being and are organized. It is worth remarking that this was so far from the interests then popular in economic theory that for more than a year I was unable to obtain a time in which to speak on the subject of Chapter 7 in MIT's own theory workshop.[5] When Tsai, with some difficulty, finally got to speak at a theory lunch, the students found the topic quite an eye-opener.

The papers reprinted in Part IB ("Associated Models of Stability Analysis") are generally earlier than those in Part IA and hark back to a simpler set of considerations in stability theory. In part because of the prominence of the gross substitutes property in the early days of tâtonnement and in part because of their role in Leontief systems, the properties of nonnegative square matrices were of considerable interest when I was a student some forty years ago. I was self-taught in linear algebra but went on to teach that subject to economists for a long time. Nonnegative square matrices provided an interesting and relevant exercise.

The theorem on such matrices that I found most natural and appealing was one that my colleague, Robert Solow, had published some years before I came to MIT (Solow, 1952). It showed that the largest (Frobenius) eigenvalue of such matrices lay between the greatest and the least of the column sums thereof (and strictly between if the matrix was indecomposable[6]). When Albert Ando and I wrote a paper for the *American Political Science Review* (Fisher and Ando, 1962), we needed to find a way to describe Solow's Theorem to political scientists (who, in those days, at least, could not be expected to know about eigenvalues and Frobenius' theorems). The result was Chapter 8, which shows that the Frobenius root is actually a weighted average of the column sums.

Solow's Theorem seemed so appealing that I wondered whether having all the column sums less than unity might not be necessary as well as sufficient for the Frobenius root to be less than one. Of course, since the column sums are not independent of the units in which the underlying variables are measured, this could not be directly true, but some thought produced Chapter 9, in which I show that it is necessary that there exist some choice of units for the underlying variables in which the

[5] Harvard was more flexible.

[6] There was a confusion of language. What many economists tended to call "decomposable" was what mathematicians tended to call "reducible" – the property that identical renumbering of rows and columns would make the matrix block triangular. What mathematicians tended to call "decomposable," the same economists tended to call "completely decomposable" – the property that identical renumbering of rows and columns would make the matrix block diagonal. I use the economists' terminology.

condition on the column sums is satisfied. Chapters 8 and 9 were written in the early 1960s, before I became really interested in stability analysis.

Chapters 10 and 11 were written in 1971 immediately after my decision to enter the subject. They were written during, and as a direct result of, Frank Hahn's visit to MIT in 1971. Hahn gave a graduate course on general equilibrium, part of which concerned stability. (Arrow and Hahn, 1971, was just being published.) Both Daniel McFadden (then also visiting MIT) and I attended the course. In one session, Hahn asked McFadden to present the paper he had written for the *Festschrift* for Sir John Hicks (McFadden, 1968). That paper showed how Hicks' remarks on stability in the appendix to *Value and Capital* (Hicks, 1939) could be justified by a model in which different markets had widely different speeds of adjustment. While McFadden's *tour de force* proved global results, its local version, at least, was based on a theorem related to those on non-negative square matrices.[7] McFadden stated that this theorem (the Fisher–Fuller Theorem[8]) was quite difficult to prove. Since it seemed to me to be of the same class of theorems as those I regularly dealt with in my class, I incautiously and immodestly suggested that it couldn't be very difficult. Fortunately, I was able to sketch the fairly simple proof, given in Chapter 10, by the end of McFadden's lecture, showing (perhaps) that I was immodest but not foolhardy.

Chapter 11 had a similar origin. Hahn happened to mention in class that he knew of no examples other than the Cobb–Douglas of a utility function generating demands that had the gross substitute property. He added that perhaps that was the only one. As someone who lectured on consumer theory and used the linear expenditure system produced by a generalized Cobb–Douglas utility function as an example, I observed that this was almost certainly not true. I then set to work to characterize the set of utility functions with the appropriate property. Because at that point I was becoming convinced by McFadden, W. M. Gorman, and others that consumer theory is best done by means of the expenditure function rather than directly by use of the utility function, I produced Chapter 11 (but not by the end of Hahn's lecture).

The chapters in Part II ("Welfare Economics and Consumer Theory") and those in Part III ("Applications of Microeconomic Theory") reflect topics in pure and applied microeconomics that have caught my attention at various times. The first of these chapters, Chapter 12, was my first published paper. It was written when I was a junior at Harvard College

[7] So-called Hicksian matrices can be written in the form A–sI, where A is a non-negative matrix, s is a scalar, and I is the unit matrix.
[8] M. E. Fisher is not related to me. See Fisher and Fuller (1958).

and being given a quite unusual education in economics by my tutor, Carl Kaysen.[9] He posed the question of what to make of the Kaldor–Hicks–Scitovsky welfare criterion if the gainers from a change did not actually compensate the losers. This prompted me to think about value judgments on income distribution; while the paper I wrote did not answer Kaysen's question, it seemed interesting in its own right.

At the time the paper was written, in the mid-1950s, welfare economics was a subject whose principal results were all negative. The discipline had emerged from a long history in which it had seemed possible to prove positive results about a normative subject, and economists had finally come largely to understand the difference between natural assumptions about behavior and value judgments. But writers were so careful not to confuse these two things, they often refused to make value judgments at all where those judgments would be likely to command less than nearly universal agreement. In particular, whenever a proposed economic measure would bring about a change in income distribution, the usual conclusion was that since one would have to make a value judgment as to whether the distributional change was good or bad, one could say no more about the problem.

In Chapter 12, I attempted to move beyond this, not by imposing my own (no doubt entirely persuasive) views as to income distribution, but by axiomatizing the properties that value judgments were likely to have. (For example, if one moves from a worse distribution to a better one, it may be natural to assume that the distributions passed through on the way are an improvement on the starting point.) Proceeding in this manner, I found it possible to reach at least some conclusions.

While I retain a fondness for Chapter 12, it certainly bore (and bears) a number of earmarks of having been an undergraduate paper written by an inexperienced author.[10] In particular, so new was I to the subject that it never occurred to me that the fact that indifference curves could not cross meant anything other than that they were parallel along rays – a very special case, indeed (that of homotheticity). For the most part, this did not matter, since the body of the paper dealt with a single-

[9] For a description of this, see the Epilogue.

[10] Aside from substance and style, it turned out to have a very high number of typographical errors, largely because it represented my first experience at proofreading. I proofread the paper with the assistance of my father, who was very helpful but persuaded me to allow him to read the galleys while I read the original manuscript. Particularly because my father was totally unfamiliar with mathematical notation, this was not a good arrangement, and I have been careful to avoid it ever since (even with mathematically sophisticated assistants).

commodity world, but it did matter for the generalization given in the Appendix.[11] This fact was picked up by Peter Kenen (then a Harvard graduate student), who wrote a critical comment. Edward Chamberlin, the editor of the *Quarterly Journal of Economics* (in which the original paper had been published), suggested that Kenen and I write jointly on the issue, and the result was Chapter 13, which showed that the slip was not of much consequence.

While I was embarrassed by the slip, I was, of course, quite excited by the adventure of publishing a professional paper.[12] Indeed, I expected the world to sit up and take notice. This did not happen – the paper did not become widely noticed and cited – but, in retrospect, I hardly feel disappointed. The paper brought me to the attention of Robert Solow (who read it before it was published) and hence indirectly to the MIT faculty. Further, the one person who did claim to be influenced by it – and who has repeatedly (and overgenerously) continued to cite it ever since in his own work axiomatizing value judgments – was Amartya Sen. (See, in particular, Sen, 1997.) If my paper was to influence only one person, that was surely the right person to influence.

My interest in the axiomatization of value judgments concerning income distribution did not end with Chapters 12 and 13. While I was at the University of Chicago in 1959–60, Jerome Rothenberg and I discussed a paper written by Robert Strotz (Strotz, 1958), with whom both of us were friendly. In a paper subtitled "A Paradox in Distributive Ethics," Strotz had adapted a set of axioms on decisions under uncertainty due to Herman Chernoff (1954) and turned them into axioms about income distribution value judgments. Strotz showed that while those axioms appeared sweetly reasonable, they had the surprising implication that all that could matter in comparing two situations was the total amount of income. Distributional considerations would cease to matter altogether.

Rothenberg and I thought this deserved a serious answer, and the result was Chapter 14, in which we argued that the apparent sweet reasonableness of the Strotz axioms was not real. Strotz replied (Strotz,

[11] That generalization was the source of some amusement in my family, all of whom were proud that I had published such a paper but none of whom were technically equipped to understand it. My aunt, Ethel Fisher Korn, was particularly amused by the last sentence of the Appendix, which, after giving the matrix generalizations of some of the theorems expressed in scalars in the text, stated, "which was to be expected." "Ah, yes," said Aunt Ethel, "I certainly expected exactly that."

[12] I was only partially gently deflated by my roommate, Richard Friedberg, who, when I proudly showed him the galley proofs, remarked, "Frank, it looks just like one of the real ones."

1961), and we added a rejoinder (Fisher and Rothenberg, 1962), neither of which is reprinted here.[13]

Chapter 14 marked my last foray into welfare economics for a considerable time. The late John McGowan and I dipped into the subject briefly when we wrote Chapter 15. That chapter deals with the question of how to decide whether advertising is excessive when it changes consumer tastes. This, in turn, involves the question of whether the consumer obtains increased utility from the consumption of advertised products.

Modeling changes in consumer taste had already been a serious concern of mine. In the late 1960s, Karl Shell and I began an investigation into the economic theory of price indices that eventually resulted in two books on the subject (Fisher and Shell, 1972, 1998). In the paper that eventually became the first essay in our 1972 book (Fisher and Shell, 1968), we considered the effect of taste changes on the cost-of-living index. We showed there that the theory of the cost-of-living index could be reformulated to accommodate such changes. In so doing, we pointed out that there is no basis in consumer theory for asserting that if someone has unchanging tastes and happens to be on the same indifference curve in two periods, then he or she is equally well off in both periods. This is because the utility levels associated with the indifference curves are arbitrary and cannot be assumed to be the same in different periods. Such an assumption has no meaning.

In the course of that discussion, Shell and I also mentioned that there was certainly no justification for the proposition that if two *different* people happen to have the same tastes and are on the same indifference curve, then they are equally happy. Yet, in effect, the notion that such a statement is "natural" persists.[14] In the mid-1980s, a series of papers by Dale Jorgenson and Daniel Slesnick followed up on a suggestion by John Muellbauer to use the model of household equivalence scales (introduced by Anton Barten as a positive descriptive device[15]) as a tool with which to make welfare comparisons. In effect, this means assuming that consumers differ in their tastes only because of parameters such as family size, age, education, and so forth, so that their utility functions can all be represented as the same function, provided one includes such para-

[13] Rothenberg, who has a legendary fondness for puns, somehow managed to persuade Strotz (against his better judgment) to give his paper the subtitle, "Paradox Regained," to go with our subtitle of "Paradox Lost." We then subtitled our rejoinder "Paradox Enow," which certainly made Strotz regret going along with us. (As the title of Chapter 7 exemplifies, I have not given up this sort of thing.)

[14] See Sen, 1997, p. 208n.

[15] See Chapter 16 for references.