Economies with Interacting Agents

Alan P. Kirman

SFI WORKING PAPER: 1994-05-030

SFI Working Papers contain accounts of scientific work of the author(s) and do not necessarily represent the views of the Santa Fe Institute. We accept papers intended for publication in peer-reviewed journals or proceedings volumes, but not papers that have already appeared in print. Except for papers by our external faculty, papers must be based on work done at SFI, inspired by an invited visit to or collaboration at SFI, or funded by an SFI grant.

©NOTICE: This working paper is included by permission of the contributing author(s) as a means to ensure timely distribution of the scholarly and technical work on a non-commercial basis. Copyright and all rights therein are maintained by the author(s). It is understood that all persons copying this information will adhere to the terms and constraints invoked by each author's copyright. These works may be reposted only with the explicit permission of the copyright holder.

www.santafe.edu

SANTA FE INSTITUTE

Economies with Interacting Agents*

Alan P. KIRMAN

(European University Institute, Florence)

April 1994

Abstract

This paper discusses economic models in which agents interact directly with each other rather than through the price system as in the standard general equilibrium model. It is suggested that the relationship between micro and macro behaviour is very different than that in the standard model and that the aggregate phenomena that can arise are rich. The models considered include ones with global interaction in which all agents can interact with each other and ones in which agents can only interact with their immediate neighbours. Both static and dynamic models are considered and the latter includes the class of evolutionary economic models.

Keywords Equilibrium, Local and Global Interaction, Communication, Dynamics, Evolution.

JEL Classification Nos C0, D0, D5.

---

* The first version of this paper was presented as the ASSET lecture at the annual ASSET conference in Barcelona. I would like to thank the participants at that conference, the participants in the IIASA conference on evolutionary game theory, and at the economics seminars at Edinburgh, Strasbourg and Berne Universities and at the Santa Fe Institute and Frank Hahn and Nick Vriend for their comments and criticism. The responsibility for what remains is still mine.

Introduction

In this paper I shall use as a starting point the most basic model in economics, the Walrasian general equilibrium model, and then go on to suggest that models in which agents interact with each other directly rather than indirectly through the market price mechanism provide a rich and promising class of alternatives which may help us to overcome some of the difficulties of the standard models. Advances made in theory in recent years, in particular those by Debreu, Mantel and Sonnenschein,(a full account of which is given in Mas-Colell [1985]) have led to increasing dissatisfaction with the Walrasian equilibrium concept, at least from a theoretical point of view. This is because the macroeconomic relationships derived from the standard assumptions on individuals are almost devoid of structure and hence of empirically testable content.

Equilibrium in the Walrasian sense can be paraphrased as a situation in which none of the individual agents has an incentive to change his choice and furthermore a feasibility constraint, that all markets clear, is satisfied. To make this slightly more precise one could rephrase this as follows. Market signals, prices, p, determine consumers' incomes and producers' profit opportunities. Consumer a chooses the best bundle of goods $\varphi(a,p)$ costing less than his income, producer b chooses that or those production plans $s(b,p)$ which maximise his profits at those prices. Equilibrium occurs when the signals p are such that

$$0 \in Z(p) = \sum_{a\in A} \varphi(a,p) + \sum_{b\in B} s(b,p) - \sum_{a\in A} e(a)$$

where A is the set of consumers, B is the set of producers and $e(a)$ are the initial resources of consumer a. $Z(p)$ is an l vector where l is the number of goods.

With appropriate assumptions on the characteristics of agents' preferences, and production technologies, the existence of such an equilibrium can be guaranteed. Although this classic notion of equilibrium is well defined and amounts to no more than finding a solution of a set of equations, it is neither uncontroversial nor very satisfactory.

The basic problem is historical. From Walras onwards, the tradition in general economics has been to maintain that such an equilibrium is one to which a market economy would gravitate. Prices would adjust in such a way as to bring about equilibrium. This notion, captured formally in the tâtonnement process, for example, and frequently referred to as the "invisible hand", conveys the idea that economic equilibrium should be thought of as the resting point of a process, and that this process should be stable. Furthermore, if analysis of the change in the equilibrium that results from a change in the parameters of the model is to be possible, one would require that equilibrium should be unique.

Some confusion has arisen from the way in which the existence of equilibrium has been proved. Given that one is looking for a solution to a system of equations of the form

$$Z(p) = 0$$

then the natural way to proceed is to define a process which modifies prices as a function of $Z(p)$ and which only fails to change prices when $Z(p) = 0$. Calling such a price modifying process $\mu(p)$ then one is looking for a "fixed point" of $\mu(p)$ either $p^* = \mu(p^*)$ if μ is a function or $p^* \in \mu(p^*)$ if μ is a correspondence. The tâtonnement process

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

is a standard example of such a process.

However the <u>existence</u> of a fixed point of such a process in no way guarantees the <u>stability</u> of the process. Indeed one only has to look at the existence proof of Debreu [1975] where he chooses an adjustment process which makes the pricing of all those goods for which excess demand is not maximal, zero. Such a process is by definition unstable but still has a fixed point.

Formal analysis has succeeded in showing under very general conditions that equilibrium exists. The particular merit of this result is seen as being that such aggregate or macro consistency is obtained from underlying "rational" or optimising behaviour at the micro level. However, to obtain stability of the equilibrium under a reasonable adjustment process or uniqueness we know that much more structure must be imposed on the system than is obtained from assumptions of the characteristics of the isolated individuals in the economy. Such assumptions can, of course, simply be on the aggregate relations by assuming , for example, that at the aggregate level all goods are gross substitutes. This sort of condition, however, cannot be derived from the standard type of assumptions on individuals. To use them would invalidate the economist's claim that aggregate analysis was based on underlying individual "rationality".

Hence, ever since the Sonnenschein-Mantel-Debreu results showed how little structure is imposed on macrobehaviour by the standard assumptions on individual economic agents, efforts to derive more macro structure from micro characteristics have increased. The standard way of circumventing this problem has a long tradition in economics and consists of considering the aggregate behaviour of the economy as though it were the behaviour of a single "representative agent".[1] This approach ensures that the problems of existence, uniqueness and stability are solved simultaneously. However, although analytically convenient, this approach is conceptually unsatisfactory and theoretically unsound (see e.g. Kirman [1992]). No doubt the length of this tradition will prevent its rapid disappearance but, recently more rigorously founded approaches have been developed which offer promising alternatives.

Remaining within the general equilibrium framework, a possible approach is to argue that it is not only assumptions on individuals'

---

[1] Reference to such a concept can be found in Edgeworth (1881) for example.

characteristics that are needed but also assumptions on the distribution of those characteristics. The representative agent approach is, of course, a special case of this, concentrating the distribution on one point. However, taking the opposite point of view, a number of authors have pursued a research programme involving the idea that heterogeneity and dispersion of agents' characteristics may lead to regularity in aggregate behaviour. This idea of focusing on the distribution of agents' characteristics with a large support can be traced back to Cournot [1843] and has been developed with considerable success by, for example, Hildenbrand [1983, 1989] and Grandmont [1987, 1992]. Hildenbrand [1994] has gone further and suggested that structure for aggregate demand can be deduced if certain regularities exist in the distribution of individuals' expenditures at different income levels. This somewhat atheoretical approach has the great merit of being empirically falsifiable.

My purpose in this paper is to go further and to suggest that we must take explicit account not only of the dispersion of agents' characteristics but also of the way in which they interact directly. This changes the nature of the link between individual and aggregate behaviour. Macro behaviour can no longer be thought of as some averaging of individual choices or outcomes and this will, as we will see, necessitate at least a re-examination of the nature of equilibrium and may even put into question the very existence of any equilibrium at all. Indeed, the relationship between micro behaviour and aggregate equilibrium, when one exists, becomes more subtle, more complex but perhaps more interesting and more realistic. Furthermore, in such models the study of the evolution of the economy over time becomes a meaningful exercise even when the state of the economy does not converge in the normal sense to any equilibrium state..

Of course, no economist has taken seriously the idea that economic agents do not interact. However, the way in which they interact in the paradigmatic model, that of general equilibrium, is very limited and very specific. As Samuelson [1963] says,

> "...individualistic atoms of the rare gas in my balloon are not isolated from the other atoms. Adam Smith, who is almost as

well known for his discussion of the division of labor and the resulting efficiency purchased at the price of interdependence, was well aware of that. What he would have stressed was that the contacts between the atoms were *organized* by the use of markets and prices."

Yet even this is not satisfactory for two important and basic features of a real economy are absent from the standard model. Firstly, who sets the prices and how, and secondly how are the trades required to clear markets, once equilibrium prices are determined, actually coordinated and achieved?

Both of these questions require for an answer some sort of communication structure which will allow for different individuals to send signals which other agents will take into account, or for a specific description of who can trade with whom. Such a structure might be deterministic or random. It might be global, in the sense that we would allow for the possibility of complete communication between all agents, or, in the random case, for uniform probability of interaction. Alternatively, it might have an explicit local structure so that one can talk about "neighbourhoods" of agents to which communication is limited.

In this paper I will examine models with this sort of structure and the sort of equilibrium notion associated with them. In addition to considering the two basic questions about the equilibrium of the standard Walrasian model that I have mentioned, such structures can be used to examine the consequences of many different types of interaction between economic agents. Individuals may be influenced by each other's preferences or choices, by each other's expectations, or only be able to trade with certain other agents. All of these types of interaction may be deterministic or stochastic as may the underlying communication network itself. Thus there could be a deterministic set of links used in a random way or the links could themselves be stochastic. Although I will divide the models that I consider into those with global and those with local interaction it will be clear that some of the formal structures used are sufficiently general to be able to incorporate both.

Global Interaction - Static models with random matching

Before examining models in which the structure of the network of links between agents is, or can be, specified I shall briefly consider those in which agents do interact directly but in which there is no specific structure determining who interacts with whom. In other words all agents can, or are equally likely to, meet each other. I will first look at those for which the equilibrium notion is static and these include many of the standard random matching models. There is a substantial literature in economics in which the role of bilateral encounters and trading has been analysed. Although no such mechanism is present or envisaged in the Arrow-Debreu model, the equilibrium concept of the latter is often invoked as a potential solution or reference point for models in which a market and trading are explicitly specified. The basic aim of Rubinstein and Wolinsky's [1990] contribution, for example, is to examine whether a process of bilateral bargaining and transaction will converge to a competitive outcome. Another earlier type of model is that of Diamond [1989] in which trading for an indivisible good takes place between many anonymous buyers and sellers. Here the equilibrium notion ceases to be necessarily a uniform price for units of the same commodity but a distribution of prices. Since agents are supposed to know the distribution of prices in the market at any point, then an equilibrium will be a distribution of prices for sellers and of reservation prices for buyers which regenerates itself. Thus one is once again reduced to looking for a fixed point of a particular transition mapping as an equilibrium.

Put simply consider the buyers in a market as being represented by the unit interval. Reservation prices are then given by a function

$$f:[0,1] \to P$$

where P is the set of possible prices and in the classic model will be a compact interval of the positive real line. Sellers are also identified with the unit interval and their prices are thus given by

$$g:[0,1] \to P$$

However f will depend on g and g will depend on f. Thus f maps a function g from [0,1] to P into another function from [0,1] to P. Similarly g maps the latter into a function from [0,1] to P. What we are looking for then are functions f^* and g^* such that

$$f(g(f^*)\cdot) = f^* \text{ and } g(f(g^*)\cdot) = g^*.$$

Of course we have to specify how f depends on g and vice versa and show the necessary continuity but it is clear that the equilibrium notion reduces to that of establishing the existence of a classic fixed point.

It is worth noting in passing that the market-clearing condition is typically dispensed with in this type of model, since sellers have an unlimited supply of the good in question at a fixed cost per unit.

Two points should be made here. In the Diamond model an individual buyer is still confronted by a market signal, the current price distribution, and responds individually to that by establishing his reservation price, whilst the seller knowing the distribution of reservation prices will set his price accordingly. Secondly, although transactions are pairwise and thus there is apparent individual interaction, the equilibrium notion is essentially a static one corresponding closely to the standard notion of equilibrium. The term "search equilibrium" conjures up an idea of active interaction, but although this type of model captures another dimension of the nature of market signals it does not really involve active interaction between economic agents. Such a criticism also holds for other models which involve random matching such as those analysed by Roth and Sotomayor [1990]. Although the solution concepts involved differ from those of the standard market in many cases, they can be sustained as competitive solutions of a market as Gale [1987] has shown. Although algorithms for finding such solutions are known, they depend in general on some central organiser and so cast little light on the market evolution which I described at the outset.

An alternative approach to studying an adjustment process which leads to a Pareto efficient outcome but which involves individual pair-wise transactions and dispenses with any central price signal, is that adopted by Feldman [1973]. He, as in the previous models, allows for random matching and requires that each pair, if they can, make a Pareto improving trade. Intuitively it might seem clear that such a process should converge to a Pareto efficient allocation. However once the notion of any central coordinating signal is removed new problems arise. Pairwise interaction or trading can easily get into a cul de sac, simply because there are possible improvements for larger groups but not for pairs. This is Edgeworth's problem of the lack of "double coincidence of wants". Coordination must be achieved in some other way. One simple solution is to require that there should be one good which every agent holds in sufficiently large quantities. This good then plays the role of money and overcomes the basic problem. However, although this type of work suggests a step forward to a genuinely decentralized mechanism it is still concerned with processes that will yield a static equilibrium of the classic sort.

Global Interaction - Dynamic stochastic evolution

I would like now to consider three classes of models which while retaining the global or uniform communication structure just described, consider the dynamic evolution of the aggregate behaviour resulting from individual interaction. In so doing, they capture some important features of the macroeconomic phenomena one would like to analyse. In each of these classes the interaction among individuals may result in quite complicated aggregate dynamics. In some cases there will be convergence to an equilibrium in the sense of a solution which will then remain unchanged over time. In others the distribution or state at any time will change continually over time and the appropriate equilibrium notion will be some sort of limit distribution of the process itself. The difference between the two sometimes rests on apparently small changes in the structure of the underlying model. The three categories of model I will treat are firstly, those which involve "herd" or "epidemic" behaviour such as those of Banerjee [1992], Bikhchandani et al. [1992], Sharfstein and Stein [1990] and Kirman [1993], and which are often thought of as being particularly applicable to financial markets. Secondly there is the type of model developed by Arthur [1989] and David [1985] to explain the adoption of new technologies when the profitability of a certain technology for a firm depends on the number of firms that have already adopted it. Thirdly there is the literature on the evolution of populations in which players are identified with strategies, are randomly matched against each other and play a game such as "prisoner's dilemma". The distribution of strategies then evolves according to their relative success (see e.g.Axelrod [1984], Young and Foster [1991], and Lindgren [1992]).

Models of "herd behaviour" suggest that there is some externality generated by the choices made by members of the population. For example in Banerjee's [1992] model individuals sequentially choose options from those indexed on a line segment. One of these options is profitable, the others are not. Each player receives a signal with probability α and this signal is correct with probability β. People choose sequentially their options. Thus observing the choices of previous individuals may reveal information about the signals that they have had. Banerjee looks for a Bayesian Nash equilibrium and finds that the equilibrium outcome will be, from a welfare point of view, inefficient, if the population is large enough, that the probability that none of the N players will choose the "current" option is bounded away from zero and that the equilibrium pattern of choices is highly volatile across different plays of the same game.

A very simple example explains the origin of the Banerjee problem. There are two restaurants A and B with an a priori probability that A is better of 51% and that B is better of 49%. However, of 100 potential clients 99 receive a signal that B is better and 1 that A is better. Thus the aggregate information suggests that B is practically certainly better. Now if the player who received the signal A plays first he will choose A provided that the probability that a signal is correct is greater than 51%. The second player observes signal B . Since signals are assumed to be of equal quality, and he knows from the first player's behaviour that he received signal A, the two signals cancel out and he chooses A using the a priori probabilities. The third player is thus left in a situation where he can infer nothing from the second player's choice, is therefore in the same position as that player was and chooses A. By the same reasoning the unsatisfactory result occurs that all clients end up at the almost certainly worse restaurant. Paradoxically both welfare loss and instability would be reduced by preventing some people from using other than their private information. Thus reducing interaction would be better for all involved. The instability that occurs when the Banerjee market is repeated is strongly related to the feedback between players and what is of particular interest is the crucial role played by the sequential structure of the moves. Bikhchandani et al. [1992] emphasise the fact that after a sufficient time the cumulated actions of other actors contain so much information that an individual will have an incentive to ignore his own information and a "cascade" will start.

In Kirman [1993] the evolution of two opinions over time in a population is discussed. The basic idea was stimulated by the observed behaviour of ants who, when faced with two apparently equally productive food sources concentrate largely on one and then focus their attention on the other. This is due to the recruiting process which is such that the more ants are being fed the stronger the trail to the source and the higher the probability of an ant leaving the nest of going to that source. Using a simple stochastic model developed jointly with Hans Föllmer it is shown that provided there is a minimal amount of "noise" in the system the proportion of ants feeding at each source will stay close to 1 or to 0 for a long time and then switch to the other extreme.

The feedback involved can either be thought of as a stronger trail or if recruiting is of the tandem type, of a higher probability of meeting a successful forager from the food source that is currently most frequented. The appropriate equilibrium notion here is then not some fixed proportion but rather a <u>limit distribution</u> of the underlying stochastic process. Thus, thinking of the state of the system as k/N, the proportion of ants at the first source, we can write $f\left(\frac{k}{N}\right)$ for the limit distribution and this should be viewed as the proportion of time that the system spends in any state k/N. Föllmer [2] showed that if one lets N become large, and approximates f by a

[2] Private communication.

continuous distribution $f(x)$ where x takes on values between 0 and 1, then this distribution will be a symmetric beta distribution, i.e. of the form

$$f(x) = x^{\alpha}(1-x)^{\alpha-1}$$

and with appropriate assumptions on the parameters of the original model the distribution will be concentrated in the tails. Thus the process will indeed spend time at each limit and occasionally switch between the two.

In a symmetric situation with a priori symmetric agents, aggregate behaviour displays violent stochastic swings.

This sort of model can be applied to a market for a financial asset as is done in Kirman [1991] and the abruptness of the change there is amplified by including into agents' observations a signal concerning which opinion is held by the majority. Here observing the behaviour of the majority provides additional information over and above that obtained by direct encounters. One has, of course, to model the market in such a way that agents do actually gain by acting with the majority.

This model is related to that developed by Topol [1991] for example and he refers to "mimetic contagion". Ellison and Fudenberg [1993a] also develop a model in which not only how well a certain choice has done but also its "popularity" is considered. They apply their model to the choice of technologies and I will turn to this in the next section.

Technological Choice

It has long been argued that as a result of externalities the value of choosing a technology may be enhanced by the fact that other firms have previously chosen it. This idea has been formalised by Arthur [1989]. He gives very simple examples of how the presence of an externality due to the number of firms that have already adopted a technology may lead a whole industry to become "locked in" to an inferior technology. Although each individual acts in his own interests the system does not generate the best collective solution. Examples of such situations are the adoption of the QWERTY keyboard for typewriters, analysed by David [1985], or the construction of light-water nuclear reactors in the U.S. in the 1950s (see Cowan [1987]).

The formal model is very simple and was developed by Arthur et al. [1983, 1984] and bears a family resemblance to the Föllmer-Kirman process mentioned earlier.

There are K technologies, x is the K vector of current shares of firms using each technology. Initial conditions are given by the vector

$y = (y_1 .... y_k)$ and the total $w = \sum_{i=1}^{k} y_i$.

Consider the vector

$$y_n = (y_{n1}, ... y_{nk})$$

which is the vector of firms attributed to each technology after $n-1$ additional units have been added to the w initial ones.

Before we can discuss the evolution of this process we must specify how the additional units are allocated. Let p define the vector of probabilities for attracting a new firm and denote by x_i the proportion of firms using technology i then we have

$p = (p_1(x), p_2(x) .... p_k(x))$

where $p_1(x)$ is the probability that a new firm will adopt technology i given the proportions that have currently adopted each of the technologies.

Note that, as Arthur [1988] remarks, one cannot invoke a standard Strong Law of Large Numbers in order to make statements about long-run proportions since the increments are not added independently. However it is clear what a long run equilibrium of this system, if it exists, should be. The mapping p takes the k simplex S^k into itself. Therefore an equilibrium will be a fixed point of this mapping.

The evolution of the process of the proportions x_n will be described by

$$x_{n+1} = x_n + \frac{1}{n+w}(b(x_n) - x_n) \qquad x_1 = \frac{y}{w}$$

where $b(x_n)$ is the jth unit vector with probability $p_j(x_n)$. Rewriting the above expression we get

$$x_{n+1} = x_n + \frac{1}{n+w}(p(x_n) - x_n) + \frac{1}{n+w}\varepsilon(x_n)$$

where ε is a random vector given by

$$\varepsilon(x_n) = b(x_n) - p(x_n)$$

The point of this is that we can decompose the system into a deterministic and a random part. Since the conditional expectation of ε given x_n is zero we can derive the expected notion of the shares as

$$E(x_{n+1}|x_n) - x_n = \frac{1}{n+w}(p(x_n) - x_n)$$

and we can think of the deterministic component of the system or the equivalent deterministic system as

$$x_{n+1} = x_n + \frac{1}{n+w}(p(x_n) - x_n)$$

Now the question is does this system converge and if so to what? Let p be continuous [3] and consider the set B of fixed points of p. Arthur [1988] gives the following results which are generalised in Arthur et al. [1987].

Theorem (Arthur [1988])

If $p{:}S^k \rightarrow S^k$ is continuous and such that the equivalent deterministic system of x possesses a Lyapounov function v whose motion is negative outside $B=\{x|p(x)=x\}$ then x_n converges with probability one to a point z in B.

Suppose p maps the interior of S^k into itself and that z is a stable point, then the process of x_n converges to z with positive probability.

If z is a non-vertex unstable point of B then x_n does not converge to z with positive probability.

The intriguing point of this analysis, which contrasts with that given in Kirman [1993] is that the industry does settle down to fixed proportions. This is due to two things. Firstly and importantly, the number of firms is continually increasing and this gives qualitatively different results from the fixed population urn model used in the Föllmer-Kirman type of process. Secondly, the restriction that the equivalent deterministic system be a gradient system is not an innocent one and rules out cycles for example. Under certain assumptions the Arthur process will actually converge to a vertex point, in which case one technology dominates the whole industry.

In modelling technological change, it is by no means clear whether the assumption of a fixed number of firms adopting different technologies or being replaced by new firms is better than that of an arbitrarily growing number of firms. Which model is chosen certainly has consequences for the nature of the results. Lastly, it should of course be said that the set of K

[3] p can actually be time variant. In other words at time n we have p_n. However, although this has the advantage of generating more initial fluctuations, Arthur et al. [1986] require that the sequence $\{p_n\}$ converge "relatively rapidly" to a limit p, at which point the analysis reduces to that given above.

technologies available should expand over time to take account of the arrival of new candidates. This, however, would add considerably to the analytical complications. Again it is worth insisting on the fact that the nature of the equilibrium depends crucially on these assumptions.

Returning to the Ellison and Fudenberg [1993a] model mentioned earlier they show that using the information given by the popularity of a technology can lead to the general adoption of the most efficient technology. The idea is simple. If agents are faced with pay-offs from one of the two technologies from which they can choose which are subject to shocks but not all agents modify their choice at each time, then watching the evolution of the choices of others reveals information about the success of those choices. If those agents who can change simply choose the best technology for the last period without looking at other firms, the result will be a limit distribution over the states of the system x_t when x is the fraction of individuals who choose technology 1. Moreover, the average of x reflects the probability in any one period that technology 1 will give a pay-off superior to that of technology 2.

Now, however, if players choose also by attaching some weight to the popularity of a technology then Ellison and Fudenberg show that if the popularity weight is chosen appropriately the system will converge with probability one to a state in which everyone uses the best technology.

In a different model, Ellison and Fudenberg [1993b] look at the consequence of sampling other agents to obtain information about their experience. Surprisingly perhaps the social learning is most efficient when communication is limited, i.e. sample size is small. Fluctuations may occur perpetually if too large samples are taken. This is, of course, related to the fact that in their earlier model there was an optimum weight to put on what amounted to a complete sampling of the population.

Once again the extent to which agents interact has a crucial bearing on the type of evolution that may occur.

Evolution in Games

A substantial literature has developed recently concerning the evolution of strategies in a large population of players who are randomly matched with each other and play a particular game. (See e.g. Hofbauer et al. [1979], Zeeman [1980] Schuster and Sigmund [1983], Bomze [1986], Hofbauer and Sigmund [1988], Foster and Young [1990] and Lindgren [1991]).

A basic notion involved in this literature is that the proportion of players playing a particular strategy should evolve according to the relative success of that strategy. An equilibrium will therefore consist either of a point in the n simplex where n is the number of available strategies or of a distribution over that simplex. Since the models involved are either based on, or invoke, biological examples, it is worth mentioning two basic differences between biological and economic models which should be kept in mind in the following discussion. Firstly, the notion of "fitness", which is the biologist's criterion for success, is not always well-defined in economic models. Secondly, human players may try to calculate the future consequences of their actions and thus the typical myopic behaviour attributed to them may not be appropriate. An a priori awareness of the goal and the desire to head in that direction may not only accelerate the evolutionary process but may also modify it. (A discussion of this sort of problem may be found in Banerjee and Weibull [1992]). However the relationship between individual rationality, however sophisticated, and fitness is tenuous, (see e.g. Samuelson [1993]).

Several approaches have been suggested to the general problem of the evolution of strategies as a function of the results from interaction with other players. One approach is to consider a deterministic rule which maps proportions of players using each strategy at time t into proportions at time $t+1$. Then one defines the "replicator" dynamics and examines the basins of attraction of the dynamic process. A second approach is to take into account the experience and information of players and to allow them to adjust their strategies individually (see e.g. Nachbar [1990], Friedman [1991], Kandori et al. [1991], Samuelson and Zhang [1992]). A last approach is to introduce stochastic elements into the replicator dynamics. Axelrod's [1984] repeated "prisoner's dilemma" tournament has been adapted for example by Foster and Young [1990], and Young and Foster [1991] to allow for uncertainty and this significantly changes the result that cooperation emerges over time.

Let me first briefly examine the case of deterministic replicator dynamics. Consider strategies in a normal form game, a large population N of players and let each of them be identified with a particular strategy. Now the outcome at each point will depend on the proportion of individuals who have chosen each strategy. In game theory, as I have said, the usual idea is that individuals meet each other at random and then play a game with the particular opponent they draw. In large populations the result of this will coincide with the expected pay-offs at the outset. Consider a matrix A, the elements a_{ij} of which give the pay-off to the player using strategy i when his opponent uses strategy j. Denoting the ith row of A by A_i and the n vector of the current proportions of the population using each strategy as p then the expected pay-off of strategy i is given by

$$s_i=A_ip.$$

Denoting by s the average expected pay-off over the n strategies then the change in the proportions is assumed to be given by

$$dp_i(t)=s_i-s.$$

One can then look for the rest points of the system and examine their stability.

This brings us to an idea introduced by Maynard Smith [1982]. He suggested that a situation could be called "evolutionarily stable" if when some small proportion ε of invading mutants enter the population they are subsequently eliminated from the population. In a two strategy game the idea of evolutionary stability has a natural interpretation. A strategy $\hat{p}$ is evolutionarily stable if $(\hat{p},\hat{p})$ is a Nash equilibrium, that is if $\hat{p}$ is a best reply to itself <u>and</u> if for any other best reply $\tilde{p}$ to $\hat{p}$, the latter must be a better reply to $\tilde{p}$ than $\tilde{p}$ itself.

As Binmore [1992] explains evolutionary stability was thought of as a useful concept since it would mean that the way in which the population evolved over time would not have to be studied in detail. Thus if we consider the n strategy situation, what can be shown is that if $\tilde{p}$ is evolutionarily stable then it is an asymptotic attractor of the replicator dynamics. Thus checking for evolutionary stability permits us to identify the attractors of the dynamics of the system. In particular when there are two players Nash equilibria play a special role. As Bomze [1986] showed, stable points in the replicator dynamics are Nash equilibria and Nash equilibria are stationary in the replicator dynamics. Yet, when more than two strategies or technologies are available there may be no evolutionarily stable situation.

However we do know something about how rational strategies have to be for them to survive, Weibull [1992] shows that any strictly dominated strategy will be eliminated along a convergent path of the replicator dynamics. Samuelson and Zhang [1992] show that even if the dynamic path does not converge, any non-rationalizable strategy will be eliminated from the population. Unfortunately, as Dekel and Scotchmer [1992] have shown, this is not true in discrete time models.

Nevertheless the idea of looking at the evolution of strategies over time in this way in a large population leads to the examination of time paths which are characterised by polymorphous populations, that is where heterogeneity persists. Furthermore it is already clear that the evolution of the aggregate system even in these simple models with random matching is sensitive to the precise specification of the model.

As I mentioned above a great deal of attention has been paid to the repeated prisoners' dilemma and much of this stems from Axelrod's famous tournament in which various strategies were matched against each other. The winner in that tournament and in a subsequent one was "tit for tat" which starts out by cooperating and then imitates the last action of the opponent. This led to the idea that cooperation would emerge as the title of Axelrod's [1984] book indicates.

However the work initiated by Foster and Young, (see Foster and Young [1990] and Young and Foster [1991]) shows how vulnerable such a conclusion is to the presence of a small amount of uncertainty. They introduce a genuinely stochastic length of playing time for each encounter and allow for mutations and immigration which always maintain the presence of all strategies.They then examine the evolutionary dynamics and define the set S of "stochastically stable equilibria". The basic idea being that the system will almost certainly be in every open set containing S as the noise goes to zero. This notion is important since, unlike the idea of an evolutionarily stable strategy, it takes into account the fact that perturbations are happening regularly and are not just isolated events. In a repeated game they show how each strategy in turn will come to dominate for a period before being superseded by the next.

All of this is done within the context of a fixed finite number of strategies. However Lindgren [1991] allows for the possible introduction of new strategies as a result of evolution. Any finite memory strategy in the repeated prisoner's dilemma framework can be written as a binary sequence which specifies the history of the actions taken by both the players and the action to be taken. This "genetic code" can then be modified by point mutations, which change one of the elements in the sequence, by split mutations which divide the sequence in half and retain one or the other half and by gene duplication which duplicates the sequence. Thus although the system starts out with strategies with a fixed memory length other strategies with different memory lengths enter the population. The evolution of the system is complicated, with periods in which strategies of increasing memory length dominate and then periods of extinction. There are also periods in which the average pay-off is low as a new mutant exploits the existing species. Such an open ended system can exhibit a number of interesting features but it is worth noting that the evolutionarily stable strategy still plays an important role.

A further line of research involving the evolution of polymorphous populations is that which has been pursued by Vriend [1994]. He constructs a very simple model of a market in which stores can sell units of a perishable good to consumers each of whom requires one unit. The price is fixed exogenously. Stores can choose how many units to stock and how many random signals to send out to consumers. Buyers can choose to buy at the store at which they were previously satisfied, to go to a store from which they received a signal or to search randomly. Each individual updates using a classifier system and in addition firms use a genetic algorithm to develop new rules. The results are striking. The distribution of firm sizes becomes rather dispersed, consumers have different behaviour and the market almost clears, 95% of consumers being satisfied. Two features are of particular interest. Firstly although the size distribution of firms is stable, individual firms do move around within that distribution. Secondly although of no intrinsic value once firms' identities are established, signals persist and even increase in quantity. This is reminiscent of the role of "cheap talk" in games.

The important contribution of this paper is that it shows how asymmetric behaviour can result from an initially perfectly symmetric situation. Although not game theoretic in the normal sense once again one can think in terms of the persistence of a variety of strategies. It has one additional feature which is that although the interaction is global at the outset it becomes local as certain consumers become tied to certain firms. This leads naturally to the next class of models.

Local Interaction: Static Models.

Here I shall look at models in which the links between specific agents and a notion of "nearness" can be defined. Consider the following question which can be posed without straying from the most basic Arrow-Debreu framework. In the standard general equilibrium model what happens if agents' characteristics are not fully determined that is, there is some randomness in them, and in particular they depend on the "environment", i.e. on the characteristics of the other agents in the economy? If the randomness affecting the individuals exhibits some sort of independence across agents, then as Hildenbrand [1971] shows, a law of large numbers can be applied and the existence of a price equilibrium can be guaranteed for a large economy.

exogenously. Stores can choose how many units to stock and how many random signals to send out to consumers. Buyers can choose to buy at the store at which they were previously satisfied, to go to a store from which they received a signal or to search randomly. Each individual updates using a classifier system and in addition firms use a genetic algorithm to develop new rules. The results are striking. The distribution of firm sizes becomes rather dispersed, consumers have different behaviour and the market almost clears, 95% of consumers being satisfied. Two features are of particular interest. Firstly although the size distribution of firms is stable, individual firms do move around within that distribution. Secondly although of no intrinsic value once firms' identities are established, signals persist and even increase in quantity. This is reminiscent of the role of "cheap talk" in games.

The important contribution of this paper is that it shows how asymmetric behaviour can result from an initially perfectly symmetric situation. Although not game theoretic in the normal sense once again one can think in terms of the persistence of a variety of strategies. It has one additional feature which is that although the interaction is global at the outset it becomes local as certain consumers become tied to certain firms. This leads naturally to the next class of models.

Local Interaction: Static Models.

Here I shall look at models in which the links between specific agents and a notion of "nearness" can be defined. Consider the following question which can be posed without straying from the most basic Arrow-Debreu framework. In the standard general equilibrium model what happens if agents' characteristics are not fully determined that is, there is some randomness in them, and in particular they depend on the "environment", i.e. on the characteristics of the other agents in the economy? If the randomness affecting the individuals exhibits some sort of independence across agents, then as Hildenbrand [1971] shows, a law of large numbers can be applied and the existence of a price equilibrium can be guaranteed for a large economy.

However, as Föllmer [1974] showed, if there is strong and complex enough local interaction between agents, then one can no longer infer the global probability law which governs the joint behaviour of all the agents from the microeconomic characteristics of the individuals. Thus Föllmer adds a new element to the standard model. He considers, for example, in an economy not only the initial characteristics of the agents $\left(\left(\succsim_a, e(a)\right)\right.$, preferences and endowments) but also a set S of possible "states", preferences and endowments of the agents and a conditional probability law $\pi_a(\cdot|\eta)$ for a in A where A is the set of agents and where η is the environment, i.e. the characteristics of the other agents.

Now if we denote by $w: A \to S$ the state of the economy and by Ω the set of all possible such states and by $\Im$ the σ field on Ω generated by the individual states, one would like to think of an economy as a measure μ on $(\Omega, \Im)$.

The problem is that once we allow for interdependence this global measure may not correspond to the individual distributions determined independently for each agent. Consistency between microeconomic characteristics and macroeconomic characteristics may be lost. Recalling that the microeconomic characteristics of the agent a are given by the probability law $\pi_a(\cdot|\eta)$then a measure μ on $(\Omega, \Im)$ is compatible with the underlying microeconomic characteristics if

$\mu[w(a) = s|\eta] = \pi_a(s|\eta)$ μ almost surely $(a \in A, s \in S)$.

Such a measure is called a macroeconomic phase of the economy and local characteristics are consistent if they admit at least one macroeconomic phase.

Various authors before Föllmer, such as Hildenbrand [1971], Malinvaud [1972] and Bhattacharya and Majumdar [1973] all studied random economies and looked for equilibrium prices when these economies were large enough. In each case some law of large numbers prevailed so that the link from micro to macro did not produce any difficulties. Without entering into details we can give the following definition for an equilibrium price vector of an economy $\mathcal{E}$ = (A,S,π) (see Föllmer [1974]). A price p equilibrates the phase μ of an economy $\mathcal{E}$ if

$$\lim \frac{1}{|A_n|} \sum_{a \in A} x(w(a), p) = 0 \quad \mu \text{ almost surely.}$$

where A_n is an increasing sequence of subsets of A which "exhausts" A. In particular, we say that p equilibrates $\mathcal{E}$ if p equilibrates every phase of $\mathcal{E}$.

In order to specify a particular type of interaction, Föllmer introduced a neighbourhood structure. That is each agent is only influenced by a limited (finite) number of other agents and this "influence structure" is defined a priori. I will come back to discuss this structure a little later. Typically, as Föllmer does, one might think of a lattice structure. In general, however, the idea is that the agent a is influenced only by his neighbours in $N(a)$ in the following sense

$$\pi_a(\cdot|\eta) = \pi_a(\cdot|\eta') \quad \text{if } \eta \text{ coincides with } \eta' \text{ on } N(a).$$

Now consider the class of economies with local characteristics which are consistent (i.e. admitting at least one macroeconomic phase) and satisfy (1). These are called Markov economies. Furthermore assume that the economy ε is homogeneous by which we mean that π is translation invariant i.e. that every agent reacts in the same way to the economy.

The importance of Föllmer's contribution is that he shows that even with this completely symmetric structure, non-homogeneities may arise at the aggregate level and it may not be possible to equilibrate markets. The point here is that no heterogeneity is required at the individual level to produce these difficulties if interaction is strong enough. Put another way, if agents are sufficiently "outer directed", that is their characteristics depend sufficiently on their neighbours, one can have the impossibility of equilibrium prices. Föllmer gives, albeit extremely stylised, examples of Ising economies in which this occurs.

Let me reiterate the essential point at this stage which is that strong local random interaction amongst agents who are <u>a priori identical</u> may prevent the existence of equilibrium. This is in stark contrast to the determinate model without this type of interaction.

It is worth noting that Föllmer worked within a static framework whilst most of the recent work on interaction and interdependence has been in the context of dynamic stochastic processes. This, however, makes his results more rather than less striking.

<u>The neighbourhood structure of an economy</u>

As soon as we introduce, as Föllmer does, a structure for the neighbourhoods of the agents in an economy, we are obliged to answer two questions. Firstly, what is the appropriate relationship to consider? Is its geographical distance as in locational models (see Gabszewicz and Thisse [1986] for example), closeness in characteristics (see Gilles and Ruys [1989] or if one is considering transactions should one not consider the probability of interaction as being related to the potential gains from bilateral trade?

In any event to discuss local interaction we must effectively impose some graph-like structure on the space of agents. One approach introduced by Kirman [1983] and developed by Kirman, Oddou and Weber [1986] and Ioannides [1990] is to consider the graph structure itself as random. Once the communication network is established, then one has to define a corresponding equilibrium notion and study the characteristics of the equilibrium that arises from the particular structure one has chosen. The difference between this and the approach adopted by Föllmer and those such as Durlauf [1990] and Benabou [1992] is that while they impose an a priori locational relationship, the stochastic graph approach can allow for agents who are "arbitrarily far" in the underlying structure to communicate. In the papers cited above, the basic idea is to consider a set of agents A and impose a probabilistic structure on the links between them. Consider p_{ab} as the probability that individual a "communicates" with agent b. In graph terms, this is the probability that an arc exists between nodes a and b. The graph is taken to be undirected, i.e. the existence of the arc *ab* implies the existence of *ba* and thus one way communication is ruled out.

In the case of a finite set A this is easy to envisage and the resulting stochastic graph can be denoted

$$\Gamma(p_{ab})$$

If there is no obvious underlying topological structure then one could consider $p_{ab} = p$, that is the probability of interaction is the same regardless of "who" or "where" the individuals are. On the other hand the structure can of course be reduced to the sort of local interaction mentioned earlier if one has an underlying "geographical" structure (see also Gabszewicz and Thisse [1986]) and one conditions p_{ab} on the proximity of a and b. The extreme case in which an integral lattice structure exists and agents communicate with their nearest neighbours is a special case of the stochastic graph with $p_{ab} = 1$ if distance $[a,b] \leq 1$.

Gilles and Ruys [1989] working within the deterministic approach developed by Myerson [1977] and Kalai, Postlewaite and Roberts [1978] adopt a more ambitious approach and use the topological structure of the underlying characteristics space to define the notion of distance. Thus agents nearer in characteristics communicate more. To an economist interested in how trading relationships are established this seems a little perverse. In general the possibility of mutually profitable trade increases with increasing difference in characteristics. Identical agents have no trading possibilities. However, there are other economic problems, the formation of unions and cartels for example, for which this approach may be suitable. The important point is that the relational structure is linked to but not identical with the topological characteristics structure.

Haller [1990] links the Gilles and Ruys approach to the stochastic one by basing communicational links directly on the topological structure of the attribute space.

The basic aim of all of these models of interaction through communication is to look at equilibria which depend on the structure of the "connected components" of the graphs involved. "Connected" here means simply that a graph Γ_A is connected if for every a and b in A there exist arcs (a,c) (c,d)... $(d.b)$ linking a and b. When looking at finite graphs the idea of the distance between a and b is simply the number of arcs in the shortest path linking the two. The <u>diameter</u> of a graph Γ_A i.e. a graph with the set of nodes A is given by

$$D(\Gamma_A) \equiv \underset{(a,b) \in A}{\mathrm{Max}} \quad \text{distance } (a,b)$$

If individuals interact with a certain probability, and these probabilities are independent, it is easy to see that the appropriate tool for analysis is the random graph $\Gamma(p_{ab})$ mentioned earlier where p_{ab} specifies the possibility of an arc between two elements a and b of the underlying set of agents A existing. Thus a particular realisation of the random graph would be a deterministic graph, undirected if a communicates with b implies b communicates with a and directed, if one is possible without the other.

To illustrate the sort of results that can be obtained let me paraphrase very informally those obtained by Kirman, Oddou and Weber [1986] and Ioannides [1990]. The former consider the problem of the core and its relation with the competitive equilibrium. Suppose that one does not allow all coalitions to form but only those that have a certain diameter, in particular diameter=1 or diameter=2. In the first case everyone in a coalition must communicate with everyone else and in the second every pair of individuals in a coalition must have a common friend in that coalition. Now consider a simple exchange economy $\mathcal{E}$ with agents A. For each realisation of the random graph on the agents A the core of $\mathcal{E}$ is well defined. Thus the core $C_i(\mathcal{E},\Gamma(p_{ab}))$ with i taking on two values 1 and 2 depending on which diameter is required is itself a random variable. If we define an appropriate "distance" between the core $C_i(\mathcal{E},\cdot)$ and the set of

Walrasian equilibrium allocations $W(\mathcal{E})$which is independent of the graph structure then a statement such as

$$p(dist.\left(C_i\left(\mathcal{E},\Gamma\left(p_{ab}\right)\right)W\left(\mathcal{E}\right)\right)\leq \mathcal{E} \quad \text{i=1, 2.}$$

is well defined.

What we would like to show is that limiting interaction or communication in the way suggested does not prevent the core from "shrinking" to the competitive equilibrium.

Consider now a sequence of expanding economies $\mathcal{E}_n$ (see Hildenbrand and Kirman [1988]). For each n let

$$p_{abn} = p_n, a, b \in A_n$$

Now Kirman Oddou and Weber [1986] obtain the following two results stated rather imprecisely

Diameter One Coalitions (strong condition).

For any $\mathcal{E} > 0$

$$\lim_{n \to \delta} \quad p\left(dist\left(C_1\left(\mathcal{E}_n, \Gamma\left(p_n\right), W\left(\mathcal{E}_n\right)\right) \leq \varepsilon\right)\right) = 1$$

if $p_n \leq \dfrac{1}{\log N}$ where $N = \# A_n$

Diameter Two Coalitions (weaker condition)

$$\lim_{n \to \delta} \quad p\left(dist\left(C_2\left(\mathcal{E}_n, \Gamma\left(p_n\right), w\left(\mathcal{E}_n\right)\right) \leq \varepsilon\right)\right) = 1$$

if $p_n \leq \frac{1}{\sqrt{N}}$ where $N = \# A_n$.

Thus, in contrast with the usual observation, it is clear that the probability of any two individuals having an acquaintance in common increases as the number of agents increases. Hence one should remark that "It is a big world" rather than the contrary. In any event with the diameter 1 we have to rely on small coalitions, but there are plenty of these and in the diameter 2 case we can use large coalitions, of which there are enough.

Ioannides' (1990) contribution is to examine the evolution of the random graph describing the communication network in his economy as the number of agents becomes large. In his model the graph becomes one with one large connected component and several smaller disjoint connected ones. He considers that Walrasian equilibria are only meaningful for economies in which the agents are linked. As a consequence a situation persists in which one price prevails in most of the economy whereas there are small islands in which the equilibria are characterised by other prices. Thus, here again, the aggregate equilibrium outcome is directly affected by the nature of the stochastic interaction between the individual agents, though here, in contrast to the previous model, connectedness is not restored as the economy becomes large.

Local Interaction: Dynamic Models.

Here I will consider some examples in which there is local interaction and in which the dynamic evolution of the aggregate behaviour and the existence of an appropriately defined equilibrium is studied.

A typical model is that of Blume [1991], who considers a countable infinity of sites, each of which is occupied by one player who is directly connected to a finite set of neighbours. Each firm can then be thought of as adopting a strategy and then receiving a payoff, depending on the strategies adopted by his neighbours. If the set of sites is δ and of strategies W, then a configuration will be $\phi : S \to W$ and the payoff to s of choosing w when his neighbours choose according to ϕ can be written $G_s(w \phi(V_s))$ where V_s is the neighbourhood of s, i.e. all sites with distance less than k from s. A stochastic revision process would then be defined such as

$$\log \frac{p_s(v \mid \phi(V_s))}{p_s(w \mid \phi(V_s))} = \beta\, G_s(v, \phi(V_s)) - G_s(w, \phi(V_s))$$

hence $$p_s(v \mid \phi(V_s)) = \left(\sum_{w \in W} \exp \beta\, G_s(w, \phi(V_s)) - G_s(v, \phi(V_s)) \right)^{-1}$$

This stochastic strategy revision process is a continuous-time Markov process. Without entering into the details, the problem is to look at the limit behaviour of this process. If β is large, that is individuals are very sensitive to the behaviour of their opponents, the process may settle down to any one of a number of equilibrium configurations. However if β is small, then the process is ergodic, that is there is an invariant measure which describes the limit behaviour of the process. Once again we see the distinction between a process which will converge to one of a number of particular "equilibrium" configurations, that is where the structure of the population is stabilised at certain distributions of strategies and that of a process which wanders through different configurations and which has, as an equilibrium, a distribution over states.

The key issue is how responsive is the probability of the choice of strategy to that of the neighbours or rather to the increase in pay-off to be obtained by changing strategy given the choice of the neighbours.

The limit when $\beta \to \infty$ is the “best response process”. In this case, if there is a Nash configuration, that is a configuration in which every firm chooses the best technology or randomises choice over best technologies if they are not unique, then one would like the concentrations of all weight on each of the points corresponding to Nash configurations to be invariant measures of stochastic revision processes. This is unfortunately not quite true. However, if in the Nash configuration the firm's best response is unique, this will be the case.

Blume [1991] examines the situation in which players play a game against each of their neighbours, and receive the average pay-off. However, it is easy to consider general pay-offs to players for adopting strategies given those adopted by their neighbours.

The important thing here once again is to see that variety can persist either through a stable distribution of strategies which remains fixed over time or through a situation in which the system spends a fixed proportion of its time in each set of configurations.

The importance of local interaction as opposed to global uniform matching is emphasised by Ellison [1993]. He considers two polar cases. First, he takes a simple two by two coordination game and considers a situation in which all of the N players are matched with uniform probability, i.e.

$$P_{ij} = \frac{1}{N-1}$$

On the other hand if we consider the players as located around a circle then one might allow players to be matched only with their immediate neighbours, i.e.

$$P_{ij} = \begin{cases} \frac{1}{2} \text{ if } i-j \equiv \pm 1 (\text{mod } N) \\ 0 \qquad \text{otherwise} \end{cases}$$

or they could be matched with many of their $2k$ nearest neighbours, i.e.

$$P_{ij} = \begin{cases} \frac{1}{2}k \text{ if } i-j \equiv \pm 1 (\text{mod } N) \\ 0 \qquad \text{otherwise} \end{cases}$$

Many intermediate cases could be constructed with probability of matching directly depending on distance on a lattice for example.

This idea, in effect, can be thought of in terms of the stochastic graph notion mentioned earlier. The rule for choice of strategy adopted by Ellison [1991] is very similar to that used by Kandori, Mailath, and Rob [1991]. A player chooses that strategy that would do best against those employed by his possible opponents in the previous period. To look at what happens consider the simple game used by Ellison

	A	B
A	2,2	0,0
B	0,0	1,1

In this game it is easy to see that a player will play strategy A if at least 1/3 of his possible opponents did so in the last period. There are two steady states (a state is the number of people playing strategy A) 0, and N. Both have large basins of attraction and although, if noise is introduced, the system will eventually move from one steady state to the other; it will do so very infrequently.

Now, however, if we look at local interaction in the same game and consider a situation in which players may be matched with any of their nearest eight neighbours, it is clear that players will choose A if at least three of these neighbours did so last time. The point here is that if there is a small cluster of players playing A it will rapidly expand and take over the whole population. Consider the case when the whole group is playing B except for players 1 through 4 who play A. Thus we have

$$(A, A, A, A, B, \ldots B).$$

Clearly at time $t+1$ players 5 and 6 will change to A as will players N and $N-1$. This experience will continue until the state with A for all players is reached. The important observation here is that if a small amount of noise (i.e., a small probability of self change) is added then it is sufficient that four adjacent players become A for the whole system to drift towards A.

This is much more likely than the $\frac{n-1}{3}$ simultaneous mutations that would be necessary for the uniform matching model to shift from all B to the basin of attraction of A.

Perhaps counterintuitively, convergence to the equilibrium, that is to the steady state distribution which puts even less weight on all B in the local matching case than in the uniform matching case, is much more rapid in the local matching situation.

Here what we see is that, in a certain sense, local interaction produces a surprisingly regular aggregate situation surprisingly rapidly.

In a related model, Ellison and Fudenberg [1993a] consider a situation in which individuals vary according to a parameter. They might vary in their geographical location or some basic characteristics. Players now learn from those with parameters close to their own. What is shown is that the system settles to a situation around a unique steady state "cut off value" of the parameter. Those with values above this use one technology, those with values below it use the other. There is a limit distribution but its variance depends on the interval of parameters over which agents make their observations. Small intervals mean low variance but clearly also show convergence. Thus a more stable limit situation is offset by a slower movement to that situation.

Another contribution in this vein is that of Bala and Goyal [1993] who consider agents who learn from those with whom they are connected. If connectivity is high enough and individuals have identical preferences then, in the long run, every agent has the same utility. However this may be a result of their having conformed to a sub-optimal action. When preferences are heterogenous, however, society may not conform at all. This depends crucially on who is connected to whom.

In concluding, it is worth noticing that in some situations learning from local experience may not be socially efficient. Convergence to a uniform situation may not occur and "stratification" may persist (see e.g.

Durlauf [1993] and Anderlini and Ianni [1993].[4] Yet in all of this the connections between individuals are given.

Suppose for example,that people in a neighbourhood who become wealthy because of their education then move out, observing those who remain would give wrong information about the value of investing in education[5]. This suggests that the notion of who is a neighbour depends on experience and this means, in terms of the earlier discussion, that the graph representing neighbourhood structures evolves itself over time.

Evolving Networks

It seems clear that the obvious way in which to proceed is to specify models in which the links between agents are reinforced over time by the gain derived from those links. Thus longstanding economic relationships would be derived endogenously from the agents experience.Instead of thinking of only the individuals learning, one could also think of the economy as learning and the graph representing the economy as evolving[6]. In the paper by Vriend [1994], discussed earlier, relationships between traders do evolve over time and a number of stable bilateral arrangements emerge. Stanley et al. [1994] consider a repeated prisoner's dilemma model in which players can choose and refuse to play against other players on the basis of updated expected payoffs. This can lead to the emergence and persistence of multiple pay-off bands.

A simple way to model an evolving graph for a fixed population of size n is to consider the set of all possible graphs, that is the 2^{n^2} nxn incidence matrices and to define a probability distribution over them. Thus a random directed graph is nothing other than a point in the unit simplex S in R^k where $k=2^{n^2}$ with the appropriate reduction in dimension for an undirected graph since the matrix is then symmetric. The evolution of the

---

[4]For results in the other direction see An and Kiefer [1993].

[5] This observation on Ellison and Fudenberg's results was made by R. Benabou.

[6]For a discussion of the formal problem of evolving networks see Weisbuch [1990].

random graph is then described by a mapping from S into S and the dynamics will then be determined by the particular form of learning used to update the probabilities attached to the links. A vertex of the simplex corresponds to a deterministic network whilst the barycentre corresponds to the uniform matching model. Careful specification of the updating mechanism should lead to interesting and potentially testable conclusions about the form of the resulting network. Thus one should be able to observe the evolution of trading groups and partnerships in markets and the development of groups playing certain strategies amongst themselves in repeated game models for example. An obvious extension of this is to consider agents as having several different types of functional links, they might be linked within a firm, as trading partners, or as members of a household for example. However the analysis of this sort of multi-layered graph seems to be much less tractable.

Conclusion

The basic argument of this paper is a very simple one. By incorporating a consideration of how agents interact into our models we not only make them more realistic but we also enrich the types of aggregate behaviour that can occur. However, as soon as we introduce this sort of interaction the notion of equilibrium has to be reconsidered, and this is particularly true if we allow for stochastic interaction and study the ensuing dynamics. The relationship between the evolution of the system as a whole and that of its micro-components may be both more complicated and different in nature to the type of aggregation with which economists are familiar. Indeed there may be no convergence at all, in any standard sense, and one is faced with analysing a constantly evolving open ended system. Nevertheless, the sort of work discussed in this paper may represent a step in the direction of thinking of the evolution of the economy as that of a self-organising system,[7] rather than as that of a glorified inter-temporally maximising individual.

---

[7]See, e.g. Lesourne (1992).

REFERENCES

An M.Y. and N. M. Kiefer, "Evolution and Equilibria Selection of Repeated Lattice Games", mimeo, Dept. of Economics, Cornell University, Ithaca, NY.

Arthur W.B. (1983), "Competing technologies and lock-in by historical events: the dynamics of allocation under increasing returns", I.I.A.S.A. Paper WP-83-90, Laxenburg, Austria. Revised as C.E.P.R. Paper 43, Stanford University.

Arthur W.B. (1986), "Industry location and the importance of history", C.E.P.R. Paper 84, Stanford University.

Arthur W.B. (1988), "Self-reinforcing mechanisms in economics", in P.W. Anderson, K.J. Arrow and D. Pines (eds.), The Economy as an Evolving Complex System, SFI Studies in the Sciences of Complexity, Addison-Wesley Publishing Company.

Arthur W.B. (1989), "Competing technologies, increasing returns and lock-in by historical events", Economic Journal, IC, pp. 116-31.

Arthur W.B., Y. Ermoliev and Y.Kaniovski (1983), "Generalised urn-problem and its applications", Kibernetika, vol. 1, pp. 49-56.

Arthur W.B., Y. Ermoliev and Y.Kaniovski (1984), "Strong laws for a class of path-dependent stochastic processes, with applications", in Arkin, Shiryayev, Wets (eds.), Proceedings of the Conference on Stochastic Optimisation, Kiev, 1984, Springer Lecture Notes in Control and Information Sciences.

Arthur W. B., Y. Ermoliev and Y. Kaniovski (1987), "Path-dependent processes and the emergence of macro-structure", European Journal of Operational Research, 30, pp. 294-303.

Axelrod R. (1984), The Evolution of Cooperation, New York, Basic Books.

Bala V. and S. Goyal. (1993), "Learning from Neighbors", mimeo. McGill University, Montreal.

Banerjee A. (1992), "A simple model of herd behaviour", Quarterly Journal of Economics, vol. 108, pp. 797-817.

Banerjee A. and J.W. Weibull (1992), "Evolution and rationality: some recent game-theoretic results", mimeo, Institute for International Economic Studies, Stockholm University, Stockholm.

Benabou R. (1992), "Heterogeneity, Stratification and Growth", mimeo MIT, Cambridge, MASS.

Brock W.A. (1993), Pathways to Randomness in the Economy: Emergent Nonlinearity and Chaos in Economics and Finance", Working Paper 93-02-006, Santa Fe Institute, Santa Fe, NM.

Durlauf S. (1990), "Locally Interacting Systems, Coordination Failure and the Behaviour of Aggregate Activity", mimeo, Stanford University, Stanford, CA.

Bhattacharya R.N. and M. Majumdar (1973), "Random exchange economies", Journal of Economic Theory, vol. 6, pp. 36-67.

Bikhchandani S., Hirschleifer D. and Welch I. (1992), "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades", Journal of Political Economy, vol. 100. pp. 992-1026.

Binmore K. (1992), Fun and Games: A Text on Game Theory, Lexington, Massachusetts, D.C. Heath and Co.

Blume L. (1991) "The Statistical Mechanics of Social Interaction", Mimeo, Cornell University, Ithaca, N.Y.

Bomze I.M. (1986), "Non-cooperative two-person games in Biology: a classification", International Journal of Game Theory, vol. 15, pp. 31-57.

Cournot A.A. (1843), Exposition de la Théorie des Chances et des Probabilités, Hachette, Paris.

Cowan R. (1987), "Backing the wrong horse: sequential technology choice under increasing returns", Ph.D. dissertation, Stanford University.

David P. (1985), "Clio and the economics of QWERTY", American Economic Review Proceedings, 75, pp. 332-37.

Debreu G. (1975), "Four aspects of the mathematical theory of economic equilibrium", Working Papers in Economic Theory and Econometrics, C.R.M.S., Institute of Business and Economic Research, University of California Berkeley.

Dekel E. and S. Scotchmer (1992), "On the evolution of optimizing behavior", forthcoming Journal of Economic Theory.

Diamond P. (1982), "Aggregate demand in search equilibrium", Journal of Political Economy, vol. XV, pp. 881-94.

Diamond P. (1989), "Search theory", in The New Palgrave: A Dictionary of Economics, (J. Eatwell, M. Milgate and P. Newman eds.), Macmillan, London, pp. 273-79.

Edgeworth F.Y. (1881), Mathematical Psychics, London, Kegan Paul.

Ellison G. (1993) "Learning, Local Interaction and Coordination", Econometrica, Forthcoming.

Ellison G. and D. Fudenberg (1993a), "Rules of thumb for social learning", Journal of Political Economy, vol. 101, no, 41, pp. 612-43.

Ellison G. and D. Fudenberg (1993b), "Word of mouth communication and social learning", Harvard Institute of Economic Research, Economic Theory Discussion Paper Series, 12, September.

Feldman A. (1973), "Bilateral trading processes, pairwise optimality and Pareto optimality", Review of Economic Studies, vol. 40, pp. 463-73.

Föllmer H. (1974), "Random economies with many interacting agents", Journal of Mathematical Economics, vol. 1, 1, March, pp. 51-62.

Foster D. and P. Young (1990), "Stochastic evolutionary game dynamics", Theoretical Population Biology, vol. 38, 219-32.

Friedman D. (1991), "Evolutionary games in economics", Econometrica, vol. 59, pp. 637-66.

Gabszewicz J.J. and J.F. Thisse (1986), 'Spatial competition and the location of firms", in Location Theory, Fundamentals of Pure and Applied Economics, Harwood Academic Publishers.

Gale D. (1987), "Price setting and competition in a simple duopoly model", Universities of Pittsburgh and Pennsylvania, first draft.

Gilles R.P. and P.H.M. Ruys (1989), "Relational constraints in coalition formation", Department of Economics Research Memorandum, FEW 371, Tilburg University.

Grandmont J.-M. (1987), "Distributions of preferences and the 'Law of Demand'", Econometrica, vol. 55, 1, pp. 155-61.

Grandmont J-M. (1992), "Transformations of the Commodity Space, Behavioural Heterogeneity, and the Aggregation Problem". Journal of Economic Theory, vol. 57, pp. 1-35.

Haller H. (1990), “Large random graphs in pseudo-metric spaces”, Mathematical Social Sciences, vol. 20, pp. 147-64.

Hildenbrand W. (1971), "On random preferences and equilibrium analysis", Journal of Economic Theory, vol. 3, pp. 414-29.

Hildenbrand W. (1989), "Facts and ideas in microeconomic theory", European Economic Review, vol. 33, pp. 251-76.

Hildenbrand W. (1993) The Law of Demand , (forthcoming Princeton University Press).

Hildenbrand W. and A.P. Kirman (1988), Equilibrium Analysis, North Holland Publishing Co.

Hofbauer J., P. Schuster and K. Sigmund (1979), "A note on evolutionary stable strategies and game dynamics", Journal of Theoretical Biology, vol. 81, pp. 609-12.

Hofbauer J. and K. Sigmund (1988), "The theory of evolution and dynamical systems", London Mathematical Society Students' Texts, vol. 7, Cambridge University Press, Cambridge.

Ioannides Y.M. (1990), “Trading uncertainty and market form”, International Economic Review, vol. 31, 3, August, pp. 619-38.

Kalai E., A. Postlewaite and J. Roberts (1978), "Barriers to trade and disadvantageous middlemen: nonmonotonicity of the core", Journal of Economic Theory, vol. 19, pp. 200-209.

Kandori M., G. Mailath and R. Rob (1991), "Learning, mutation and long run equilibria in games", mimeo, University of Pennsylvania.

Kirman A.P. (1983), "Communication in markets: a suggested approach", Economics Letters, vol. 12, no. 1, pp. 101-8.

Kirman A.P. (1991), "Epidemics of opinion and speculative bubbles in financial markets", chapter 17 (pp. 354-68) of Money and Financial Markets, M. Taylor (ed.), Macmillan, London.

Kirman A.P. (1992), "Whom or what does the representative individual represent?", Journal of Economic Perspectives, vol. 6, 2, pp. 117-36.

Kirman, A.P. (1993), "Ants, rationality and recruitment", forthcoming, Quarterly Journal of Economics, February.

Kirman A.P., C. Oddou and S. Weber (1986), "Stochastic communication and coalition formation", Econometrica, January.

Lesourne J. (1992), The Economics of Order and Disorder, Clarendon Press, Oxford.

Lindgren K. (1991), "Evolutionary Phenomena in Simple Dynamics", in Artificial Life II, C.G.Langton, C. Taylor, J.D. Farmer, & S. Rasmussen (eds), Addison-Wesley, Redwood City, CA.

Malinvaud E. (1972), "The allocation of individual risks in large markets", Journal of Economic Theory, vol. 4, pp. 312-28.

Mas-Colell A. (1985),The Theory of General Economic Equilibrium: A Differentiable Approach, Cambridge University Press.

Myerson R.B. (1977), "Graphs and cooperation in games", Mathematics of Operations Research, 2, pp. 225-9.

Nachbar J. (1990), "'Evolutionary' selection dynamics in games: convergence and limit properties", International Journal of Game Theory, vol. 19, pp. 59-89.

Roth A. and M. Sotomayor (1990), Two sided Matching: A Study in Game-Theoretic Modelling and Analysis, Cambridge University Press, New York.

Rubinstein A. and A. Wolinsky (1990), "Decentralized trading, strategic behaviour and the Walrasian outcome", Review of Economic Studies, vol. 57, pp. 63-78.

Samuelson P.A. (1963), "Problems of methodology: discussion", American Economic Review, vol. 54, pp. 231-36.

Samuelson L. and J. Zhang (1992), "Evolutionary stability in asymmetric games", Journal of Economic Theory, vol. 57, 2, pp. 363-91.

Samuelson L. (1993), "Recent Advances in Evolutionary Economics: Comments." Economics Letters, vol.42, 2/3

Schuster P. and K. Sigmund (1983), "Replicator dynamics", Journal of Theoretical Biology, vol. 100, pp. 533-8.

Sharfstein D. S. and J. C. Stein (1990), "Herd behaviour and investment", American Economic Review", vol. 80, 3, pp. 465-79.

Smith J.M. (1982), Evolution and the Theory of Games, Cambridge University Press.

Stanley E.A., D. Ashlock and L. Tesfatsion (1993), "Iterated Prisoner's Dilemma with Choice and Refusal of Partners" Economic Report no.30, Iowa State University, Ames, Iowa.

Topol R. (1991) "Bubbles and volatility of stock prices: effects of mimetic contagion", Economic Journal, vol. 101, July, pp. 786-800.

Vriend N. (1994) "Self-Organized Markets in a Decentralized Economy", Working Paper 94-03-013, Santa Fe Institute, Santa Fe, NM.

Weibull J.W. (1992), "An introduction to evolutionary game theory", Department of Economics, Stockholm University, mimeo.

Weisbuch G. (1990), Complex System Dynamics, Addison-Wesley, Redwood City, CA.

Young H.P. and D. Foster (1991), "Cooperation in the short and in the long run", Games and Economic Behavior, vol. 3, pp. 145-56.

Zeeman E.C. (1980), Population Dynamics from Game Theory, Lecture Notes in Mathematics 819, Springer Verlag.