DRAFT

1. Introduction

The papers in this series are linked by a common theme. As the frontispiece to the volumes proclaims: 'In contemporary economics the label 'theory' has been appropriated by a group that confines itself to largely a-social, a-historical, mathematical 'modelling'. *Economics as Social Theory* thus reclaims the 'theory' label, offering a platform for alternative, rigorous, but broader and more critical conceptions of theorising.'

The mere existence of a general commonality of view does not, of course, guarantee that the particular offerings will meet with universal approval. Mao Tse Tung's injunction to let a thousand flowers bloom might indeed have been meant at its face value, neither more nor less, as an encouragement to diversity. But it has often been interpreted as a way of inducing the revelation of true opinions in order that heretical ones might be more easily identified and suppressed.

It must be said from the outset that this paper risks such dissolution in a suitable hail of fire and brimstone. For it does deal with mathematical modelling. In mitigation, before the scales are weighed and judgement passed, several points must be made.

Conventional economics pays very little attention to the ways in which individuals are linked together. The aim of this paper is to show three things. First, that this is an important omission from orthodox thinking. Second, that taking account of this in a systematic way can give powerful insights into social and economic issues. Third, that formal modelling is a useful way of doing this, which helps in the process of developing alternatives to orthodoxy.

An important assumption of orthodox theory is that all agents - individuals and firms, in everyday English - have access at all times to all available information. Further, it is assumed that the tastes and preferences of individuals, whilst they differ, are nevertheless fixed. Each individual acts in isolation, maximising, in the jargon, his or her utility. The actions of others only impinge indirectly on the agent, via their impact on the set of prices which the individual faces.

The internal contradictions of this core model of economics have been explored rigorously by mathematical economists. I have described much of this work in both the *Death of Economics* and *Butterfly Economics*, and it is not my purpose to repeat this here[1]. Economists have responded to these problems in the past twenty years or so. A literature has grown in which individuals are assumed to have access to 'imperfect' information. The choice of phrase economists use to describe this theory is a telling one. For it implies that the world of free market theory is the 'perfect' one, the Platonic Idea towards which all actual economies must strive, and the messy world of reality is 'imperfect'. But even slight deviations from perfection in terms of the basic assumptions which are made can lead to outcomes and implications which are far removed from the world of perfect competition.

The first example in this paper illustrates exactly this concept, and the second deals with the importance of social networks in the process of how the unemployed find jobs. In both cases, the way in which economic agents are connected together for the exchange of information is of decisive importance for the outcome.

2. Market structures

The core world of economic theory consists of a very large number of small firms. No firm is of sufficient size to be able to exercise any control over the market for its product. An important implication of this is that companies are what is known as 'price takers'. In other words, each individual firm has to accept the market price for its product, and has no possibility of influencing it in any way.

In this make believe world, firms can sell as much output as they choose to produce. There is no real incentive to increase production, because the forces of competition have eroded profit margins to a 'normal' level. However, if a firm decides to, say, double its output, it does so under very special conditions of production. In a sentence, the costs per unit of output never change regardless of whether a firm produces just one unit of output or a hundred million units. The cost per unit is always the same.

In the jargon, this assumption is called 'constant returns to scale'. If the inputs which the firm uses in the process of production are doubled in size, output is also doubled. In short, there are no particular advantages or disadvantages to being either big or small.

In reality, of course, this is far from being the case. The desirability of certain products depends almost entirely on the fact that they cannot be produced on a large scale. At one level, we might regard all red wine as being the same in the sense that each variety is a red, alcoholic drink made from grapes. This differentiates it sharply from other alcoholic drinks such as whisky, and non-alcoholic drinks such as milk. But only the most puritanical teetotaller can fail to be aware of the myriad of distinctions within the red wine market. For connoisseurs, the pleasure of drinking a *premier grand cru* Bordeaux would surely be diminished if the wine could be made on the same scale as, say, mass market fiery Algerian or Bulgarian red.

But more generally, there are great advantages to being big. Most markets are dominated by a few, very large firms. The costs of producing a unit of output, far from being constant, often fall as the amount which is produced grows. The decidedly Old Economy industry of car production illustrates this well. A dedicated and eccentric amateur could possibly manufacture his or her car from scratch in the back yard, though this would take a very great deal of time and hence cost. For mass production, a large

---

[1] Readers interested in the formal details should go to Roy Radner's brilliant paper 'Competitive Equilibrium under Uncertainty'. Published in *Econometrica*, a world class journal, as long ago as 1968,

and expensive plant is required. Once assembled, if the plant makes just one car per year, the costs of producing it are enormous. For in addition to the small amount of additional labour costs needed to produce it, all the costs associated with building the plant fall on this one car. If two cars are made, the costs per car are already halved, as each car then shares the total costs. And so on and so on, until the limit of the capacity of the plant is reached. This state of affairs is known in economics as 'increasing returns to scale'. In other words, as the scale of output increases, the advantages of being able to produce more increase.

The car industry, indeed, contains the world's largest company, namely General Motors with an annual value of sales of around $200 billion. But many other industries are dominated by a small number of firms. The various facets of the computer industry share this characteristic, with Microsoft being the best known example. The world of advertising might be thought to be immune from such trends, based as it is on the individual skills of the people employed rather than relying on huge amounts of investment with which to produce its output. In the cynical phrase, 'the assets walk out of the door at five o'clock every night'. But even here, where there is scope for lots of small businesses to survive by creating niches, the industry is dominated by a few large firms - J Walter Thompson, Publicis, Ogilvy and Mather and the like. Huge multi-national companies often prefer to have their advertising accounts handled on a global scale by a single agency, and only the very biggest are capable of doing this.

In short, the typical market structure in the Western economies appears to be far removed from the world of perfect competition. Instead of many small firms, industries are dominated by a small number of very large ones. 'Oligopoly' rather than perfect competition is the empirical norm.

All these facts are well known. Indeed, they are an important reason why many students abandon the discipline of economics at an early stage. The world of economic theory seems remote from reality. This is a perfectly reasonable judgement to make.

---

this single paper completely undermines the theoretical basis of orthodox economics.

After all, many of the assumptions of conventional economics were made purely to make the maths tractable. When economic theory was first being formalised in the late nineteenth and early twentieth centuries, economists had to formulate their models in such a way that they yielded what are known as analytical solutions. In other words, the equations in the models could in principle be solved using a pencil and a sheet of paper. It might be difficult, but it was nevertheless feasible. In order to do this, the assumption of constant returns to scale was needed. Great economists such as Alfred Marshall, who dominated the subject around the turn of the nineteenth century, knew perfectly well that this was often not true. But it was the best which could be done at the time.

The closing years of the twentieth century have seen an explosion in computing power which enables more complex systems of equations to be solved for the first time. We need no longer rely on pencil and paper, but can obtain what are called numerical solutions - using the power of the computer to grind them out. The advantage of this is highly practical. More realistic models can be built.

How do economists justify the apparently glaring disparity between the oligopolistic structure of the real world with the perfectly competitive one of economic theory? There are two main defences. The first is the classic one of 'as if'. It may seem - indeed it may even be the case - that many industries are dominated by just few firms. But competitive pressures are such that it is 'as if' there were a large number of small firms in these industries. To be fair, this argument has a certain amount of validity. Car companies, despite their enormous size, as feeling the pinch as car ownership amongst Western households approaches saturation levels, at least as far as first cars are concerned. It is becoming more and more difficult for them to increase prices. A newcomer like Microsoft, which it was in the 1980s, can enter a market characterised for many years by the towering presence of IBM.

Yet, with the best will in the world, it *is* hard to imagine that these industries are essentially the same as the theoretical world of perfect competition. Certainly, there are many able free market economists who would long to receive the sorts of consultancy fees available to those hired in the various anti-trust actions brought against Microsoft by the US authorities. But despite the huge potential supply, the price - the fees - of those who were chosen were not eroded to a perfectly competitive level.

The second argument used by economists is not even as good as the first. Confronted by the very many examples of oligopoly, there is a general waving of hands, and a murmuring about how this arises because of increasing returns. In some way, it is the real world which is seen as the aberration. Perfect competition is the ideal norm, and all other forms of market structure are a deviation from it.

Using the power of computers, we can set up models which demonstrate that this is *not* the case[2]. Almost all of the assumptions of the perfect competition model can be retained. But relaxing just two of them, in other words, making the model more realistic, leads to outcome in which oligopoly and not perfect competition is the norm - exactly as we observe in reality.

The basic model may still seem unrealistic, but that in itself is not a criticism. All models must by definition be approximations to reality. But the purpose of this model is to show that we do not need to invoke increasing returns to bring about situations in which oligopoly is the norm. Of course, increasing returns actually do exist, and exist pervasively. So in this important sense the model is deliberately *un*realistic. It illustrates an important aspect of the *process* of how a competitive structure evolves into an oligopolistic one[3].

We start with a large number of firms operating under conditions of constant returns to scale. Just as in conventional theory, the product produced by each of these firms is identical. We might imagine these firms as being small companies which supply, say, the motor vehicle industry. More precisely, these are firms which supply firms

[2] The first example in the chapter can, in its simplest form set out here, actually be solved analytically. But this soon ceases to be the case when even slightly more complex and realistic assumptions are made.
[3] My colleagues at Volterra, Craig Mounfield, Bridget Rosewell and Laurence Smith are all actively involved in this ongoing research

which supply firms which themselves supply the likes of Ford and General Motors. They manufacture, for example, bits of metal which are used to make bits of gearboxes. We need to think of this kind of situation, in which companies are supplying fairly standard industrial parts to other companies, in order to make the assumption that their product is identical at all realistic. In many markets, firms make a great deal of effort to convince people that their product is different from that of other companies even when it is essentially the same. Air travel, for example, consists of transporting people long distances in metal tubes at high speed. The speed is its key advantage over other forms of travel. Travelling, say, from London to New York, it does not matter which airline is chosen in terms of getting there much faster than by ship. Yet airlines spend huge amounts to convince customers that they are different from each other, efforts which by and large are successful.

One of the variations we make in our assumptions from those of conventional theory is that the companies making the product are not identical to each other. They make the same thing, but with different levels of costs. We can rule out differences in the costs of transporting the product to the buyer by imaging these companies to be part of a geographic cluster, so both the buyers and the sellers are close to each other. In areas with large car plants, myriads of small and medium sized firms have developed over the years in the same places, each supplying component parts into the final product. This also rules out differences in labour costs in the area where production takes place. Cost differences arise because they produce with varying degrees of efficiency. This could arise for many reasons, for example, from the existence of different vintages of the capital stock needed to produce the commodity. Once again translating from economic jargon into English, this means that some of the firms use older and less efficient machines than others.

An enormous amount of empirical work demonstrates that firms in the same industry do differ, often substantially, in efficiency. This is a source of embarrassment, for two reasons. First, it is often hard to explain exactly why companies do differ. The marvellous concept of 'X-efficiency' has been invoked to account for differences. To a cynic, this could very well mean 'we don't know how to explain this phenomenon, but, like Rumplestiltskin, now we have given our ignorance a name we feel much better about it'. However, being charitable we could think of X-efficiency as arising from different qualities of management. All firms could have identical capital stocks, labour costs and so on, but management efficiency might vary across the firms. For our the model to operate, we only require that firms differ with respect to a single factor.

The second difficulty is more profound for orthodox economics. Competition is meant to involve the survival of the fittest. Everyone is forced either to drive costs down to those of the most efficient producer, or to go out of business. For, by definition, the price of the product cannot be affected by individual firms and so higher cost cannot be passed on into higher prices. Yet in the real world, inefficiencies can persist for very long periods of time. Study after study, government report after government report have shown, for example, that whilst the very best firms in the UK are as good as any in the world, a large proportion of British business transparently is not.

This leads into the second difference in the assumptions of our model from those of the perfectly competitive world of conventional theory. There is a large number of producers and we assume that there is an equally large number of buyers. Each of the producers is making the identical product at its own price, which differs from that of every other firm. But, and this is the second difference in assumptions, there are substantial costs of acquiring information about these prices. One practical reason for this might well be the fact that the Managing Director of the firm you have always bought from is your cousin, or your golf partner, or a member of your Masonic Lodge. There would be non-monetary social costs associated with buying from someone else.

In a more realistic setting, in which firms are geographically dispersed for example, information may be costly to acquire for many reasons. One of the advantages of the Web is that it enables businesses buying from other businesses - B2B - to acquire quotes from around the world very cheaply. Yet, paradoxically, this apparent enhancement of competition may trigger a process which leads to more and more oligopolistic market structures.

Suppose in our imaginary world that each buyer only gets a quote from a single supplier, and buys one unit of output from this seller at the quoted price. We have an equal number of buyers and sellers, and we assume that the buyers are connected to the sellers at random. There are no other connections in the model. An important practical example of this latter feature would be to specify a network of connections between the sellers, across which information about their prices could be transmitted, so that collectively they might exercise some control over the market price. Of course, active collusion on these lines is illegal in many countries, but there are many ways and means of signalling such information. In consumer markets it is very easy. Prices are visible for all to see, so companies can perfectly legally all charge the same price without any suggestion of collusion. In business to business markets it is a little harder, and anti-trust cases often turn on whether or not companies can be judged to have taken part in price-fixing agreements, explicit or otherwise.

We can now examine the effect of reducing the cost to buyers of acquiring information in this market. We have set the simple rules of the game, as it were, and can observe how it unfolds. We start off with, say 100 buyers each connected at random to a single one out of 100 sellers, so that buyer number 53 might be connected to seller number 34 and so on. Because of the way the connections are specified, some sellers will not be connected to any buyer initially. In other words, they make no sales. The precise number will differ each time the game is played, but usually between 30 and 40 out of the 100 sellers will find themselves in this position. We can introduce rules about how long such sellers can remain in the game, or we can think of them as potential rather than actual participants. To keep the rules as simple as possible, we choose the latter. Firms which make no sales do not go out of business, but are able to quote and price and produce subsequently if they are asked to do so by a buyer.

Apart from the firms with no sales, most of the sellers will start by being connected to just one buyer, so they produce one unit of output. A small number will be connected to two buyers, an even smaller number to three, and very occasionally, if we play the game lots of times, we will start with some sellers having four or five buyers. But in general, at the start of the game, the sellers are all producing broadly similar amounts of output.

The game now moves onto the next step. This time, each buyer is connected to two sellers. One of these is the previous selection, and the new one is chosen at random. We can think of this happening because the cost of acquiring information has fallen. For the same amount of cost and effort involved in dealing with one seller, buyers can now contact two. Each buyer ruthlessly selects the lowest price. There is no difference in the quality of the product, whether actual or merely perceived, which will promote buyer loyalty to the previous seller. And buyers switch supplier even if the difference in price between them is tiny. We can add features such as allowing suppliers to make a one-off reduction in their price if they would otherwise lose a sale, but this makes no difference to the eventual outcome.

At each subsequent step of the model, an additional connection between each buyer and the group of sellers is added. With each step, the extra connections can be interpreted as further reductions in the cost of obtaining information. One point to notice is that this is subject to diminishing returns. The first additional connection doubles the amount of information obtained from one quote to two, the next increases it by fifty percent, from two to three, and so on in ever diminishing proportions.

The buyer obtains a price from the previous firm with which he or she did business, and a number of other quotes chosen at random. Just as before, the lowest price is chosen without sentiment. Low-cost sellers will begin to obtain more and more of the total market. Eventually, of course, all buyers will be connected to all sellers. When this happens in our game, everyone will choose the lowest price seller. The competitive circumstances in which the game began will evolve into a monopoly, with just one seller.

As it stands, the game is a simple one, but it nevertheless informs about the process of how an apparent increase in the degree of competition eventually leads to less rather than more competition. The chart below plots a typical solution of the game showing the number of active sellers - those firms making sales in the current period - and the number of connections from each buyer.

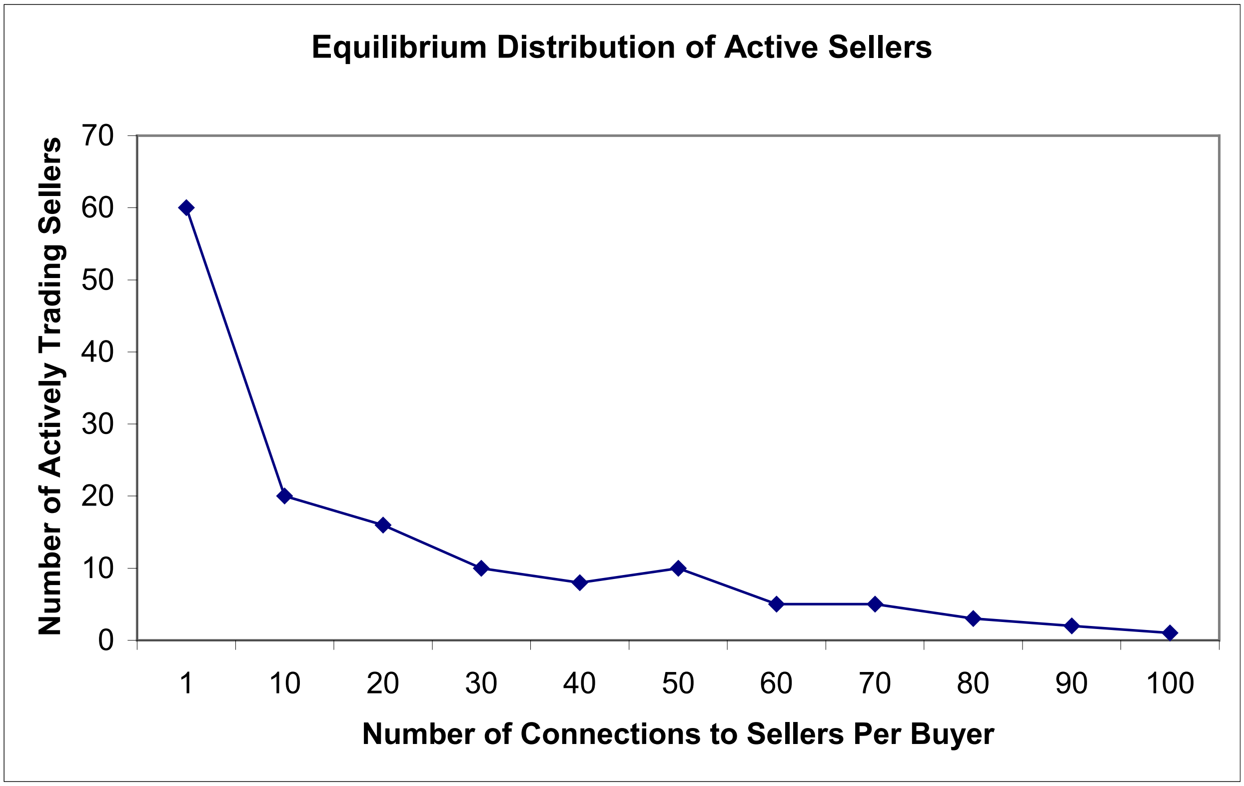

To begin with, in this particular outcome of the game, we have just over 60 sellers active in the first step - those which are connected to a buyer. The most interesting point about the chart is that the number of active sellers falls away initially very rapidly as the number of connections increases. Further, the size of those which remain active begins to vary substantially. In the above example, with just 10 (out of a possible 100) connections from each buyer, the number of active sellers falls from over 60 to 20. The largest seller is now producing 17 units of output and the smallest (ignoring those not able to sell anything at all) just one. This is qualitatively very similar to the sort of market structure which we observe in many real world markets.

The model, or game as we have been calling it, can be made more realistic. For example, we can specify networks on which the sellers are connected, as mentioned above, and allow price collusion to emerge once sellers on the same network acquire more than a certain percentage share of the total market.

But even in a very simple form, this illustrates how markets in which firms are producing a very standardised product can evolve into oligopolies. By definition in this model, firms produce under constant returns, so we do not need to invoke increasing returns to get this result. Further, all firms produce an identical product, thereby eliminating the scope for creating partial monopolies or oligopolies by product differentiation, another reason often put forward to account for the market structures which we actually see. The only departure we require from the orthodox theory of perfect competition is that firms differ in their efficiency, a pervasive feature of reality. As the costs of acquiring information fall and the number of connections rises, perfect competition gives way to oligopoly. The ideal outcome of economic theory, with many small firms producing under identical circumstances, should be seen as the particular rather than the general case. Unless *all* of the theoretical conditions required are met, oligopoly and not perfect competition is the standard outcome towards which markets evolve.

3. Unemployment[4]

A familiar and depressing feature of the developed economies is the way in which high pockets of unemployment can persist in certain areas even when the overall demand for labour is high. In the South of England at the end of the year 2000, for example, there is effectively full employment. In many areas the unemployment rate is less than 1 per cent, a rate not seen since the Golden Years of the 1950s and 1960s. Yet in particular

[4] An expanded version of this section is given in a paper I presented with Laurence Smith to the Tenth World Congress of Social Economics held at the University of Cambridge, August 2000

localities, often in close geographical proximity to very prosperous places, the unemployment rate is in double figures.

In part, this is due to the low skills and educational attainments of people in these neighbourhoods, but the phenomenon can be seen at an even more local level, comparing employment outcomes for residents of neighbouring housing estates or schemes, or often within individual schemes themselves. Despite broadly similar skills - human capital as economists like to describe it - some people on the estates appear to find it easy to get a job whilst others struggle and remain unemployed.

A key way in which people find jobs is through personal contact. Mark Granovetter, the American sociologist who did seminal work in this area in the 1970s, reported, for example, that over 60 per cent of professional and managerial workers in the US obtained their jobs through personal contact. The same phenomenon occurs much lower down the jobs hierarchy. The infamous British Old Boys' Network was often invoked as a reason why Etonians or Harrovians found themselves in top jobs, but the same sort of thing happens at all levels of the jobs market.

For example, Carmel Hannan reported in a University of Essex paper in 1999 the results of a detailed statistical study of micro-data held in the British Household Panel study, examining the impact of both traditional economic explanations of exit from unemployment such as the level of human capital of the individual, and the social network hypothesis. This latter is found to be very important. Pamela Meadows, former chief economist at the UK's Department of Employment, wrote in 2000 that 'The use of friends, relatives and social contacts as a means of getting jobs remains very important to young people and is attractive to employers. However, this works to the disadvantage of those who are outside these networks, because their parents and friends are unemployed, because they have lost contact with their family or because they have a bad reputation or criminal record'. She went on to point out that a key reason for this is the attitude of employers: ' informal methods work in a variety of ways. There are traditional introductions by fathers, uncles, brothers and family friends, who essentially take the role of sponsor to the young person, and provide some guarantee of his worth to the employer. This has always been an important route into work for young men and it remains surprisingly strong. ..... From the point of view of the employer, the grapevine is seen as less risky than recruiting in the open market, because it includes an element of social control'.

This section of the chapter sets out a highly stylised model of unemployment and job search. The aim is to focus explicitly on the nature of the social network which connects agents in the model, and on the implications for unemployment of different degrees of connectedness of social networks.

We set up a model in which unemployment can persist only because of a failure of the social network to inform individuals about job opportunities. It must be stressed that we are not suggesting that this is the only reason why unemployment exists. But to the extent that a lack of information about job vacancies might be a cause of unemployment, we quantify the influence on this of the type of social network by which individuals are connected and information exchanged between them.

An important result which emerges is that, even when jobs are being created, high rates of unemployment can emerge in social networks which are only weakly connected. And as noted above, we do observe high pockets of unemployment even in the face of a strong demand for labour. Abstract though the model is, it points clearly to the kinds of policies which are needed in order to deal with this problem.

We have a model populated by a largish number of individuals. We might think of it as representing a housing estate or other local area. At any point in time, the actors in the model are either unemployed or employed. In other words, the focus is on people of working age. In the most straightforward version, the assumption is made that everyone is identical. Patently, this is not true, but as an approximation of the sorts of skills and qualifications which might be found in a rundown inner city housing block, for example, or in one of the massive estates which surround certain British and French cities, it is not completely unreasonable.

In this model, job losses and the creation of new jobs takes place at random. In other words, the motivations of firms in reducing/expanding employment are not considered in this model, and neither is the overall economic climate in which firms operate. We could imagine a situation in which the economy as a whole is close to full employment. Even in these conditions, job losses still occur, but are balanced out by job creation. And from the perspective of the low skill inhabitants of a housing scheme, the job gains and losses in nearby companies is essentially a random process. There is little which they feel they can do to influence the outcome. They are unlikely, for example, to be able to borrow large amounts of venture capital to create the new Microsoft.

In our artificial world, if a job vacancy arises and a person hears about it, he or she will apply for the job and be hired with certainty. We can make this more realistic by giving individuals their own probabilities both of applying and of being accepted by the employer. But the simpler version illustrates the same principles.

To complete the model, or the rules of the game, we need to specify how it evolves. As with the competition model in section 2 above, the action takes place in consecutive steps. Imagine to begin with that everyone is out of work. As it happens, the eventual outcome is the same no matter how people start off, but the details of demonstrating this need not concern us here. In the first step of the model, two things happen. First, someone is chosen at random from the population. If already unemployed, he or she remains unemployed. But if the person has a job, it is deemed to be lost and the person becomes unemployed.

The second part is a bit more complicated. Another random selection is made from the people in the model. If the person chosen is unemployed, he or she is offered a job and takes it. But if the person is already employed, the process does not stop there. The job offer is passed on to those people on the individual's social network. In other words, to friends, family or neighbours - to whoever the person is likely to tell the news about a job being available. From this group, anyone who is unemployed takes up the job[5]. This is exactly the process described above in empirical accounts of an important way in which actual labour markets work. The network both informs people about job vacancies, and helps them secure it when they apply. The mode then moves on to the next step. For completeness, if everyone on the network is employed, the job offer disappears from this group completely. It is not carried over to the next period. This seems quite realistic: a firm has a vacancy and asks some of its workers if they know anyone who would be suited. If they do not, the company will start looking elsewhere.

The process then continues, step by step, and the mixture of employed to unemployed in the group of people in the model evolves. By experiment, examining many different solutions of the model on the computer, we can judge when this split between the employed and unemployed settles down. At every step there may of course be a change, but these eventually become small, and a proportion which is effectively constant emerges.

The relative size of the employed and unemployed groups which the model generates depends crucially upon how we specify the social network. We can judge the eventual outcome by reference to the two extremes. First of all, suppose we have a very gregarious, integrated community in which everyone knows everyone else and tells them about job offers. Of course, no actual group is ever like this. But with the professional groups examined by Granovetter which we discussed above, well over half the people obtained jobs through personal contacts. The outcome of the game in this situation is easy to see. In the first step, someone is bound to be offered a job because everyone is unemployed. In the second step, even if by chance this same person is selected to be offered the new job, he or she will tell everyone else, and so the job will be filled. Every time someone in work is told of a new job offer, everyone else in the population is told about it. So eventually everyone ends up employed.

[5] If more than one person on the relevant network is unemployed, one of these is selected at random to take the job

The other extreme is a very disjointed, anomic group in which people keep to themselves and do not communicate directly with anyone else. In many ways, this is the model of free market economics. No-one passes on job offers to anyone else and operates as a completely selfish individual. The outcome in this case is that the unemployment rate settles down at 50 per cent of the population: half find work, but the other half do not. It takes about a page of straightforward algebra to demonstrate this result, and it could be described in words. But this would be a lengthy diversion. The intuitive interpretation of this is that at any step of the model, anyone who is chosen has a 50/50 chance of being made unemployed (the first part of each step), or of being offered a job (the second part). If we allow the model to run for a sufficient number of steps to give everyone the chance of being chosen in each part, this is the overall split which will emerge in the population.

Imagine we have 100 people in our model. In the two cases examined above, any particular individual has a social network with either 100 or 0 connections. In the first one, everyone is connected to everyone else, and passes information on. In the second, no-one has any connections to anyone else. And the unemployment rate is either zero or 50 per cent, depending on which network we specify.

What happens when we have a different number of connections per individual? How rapidly does the unemployment rate fall with the number of connections, with the closeness of the social networks? Here, we usually do need the computer to obtain the result. There are many different ways in which people can be connected, and there is a whole branch of mathematics - graph theory - which deals with this sort of situation. The precise results do depend on the exact geometry of the connections, but the most important influence on the outcome of our model is how many connections, on average, each individual has.

The chart below illustrates how the eventual rate of unemployment which emerges depends upon the density of the social network, on how well connected and integrated people are[6].

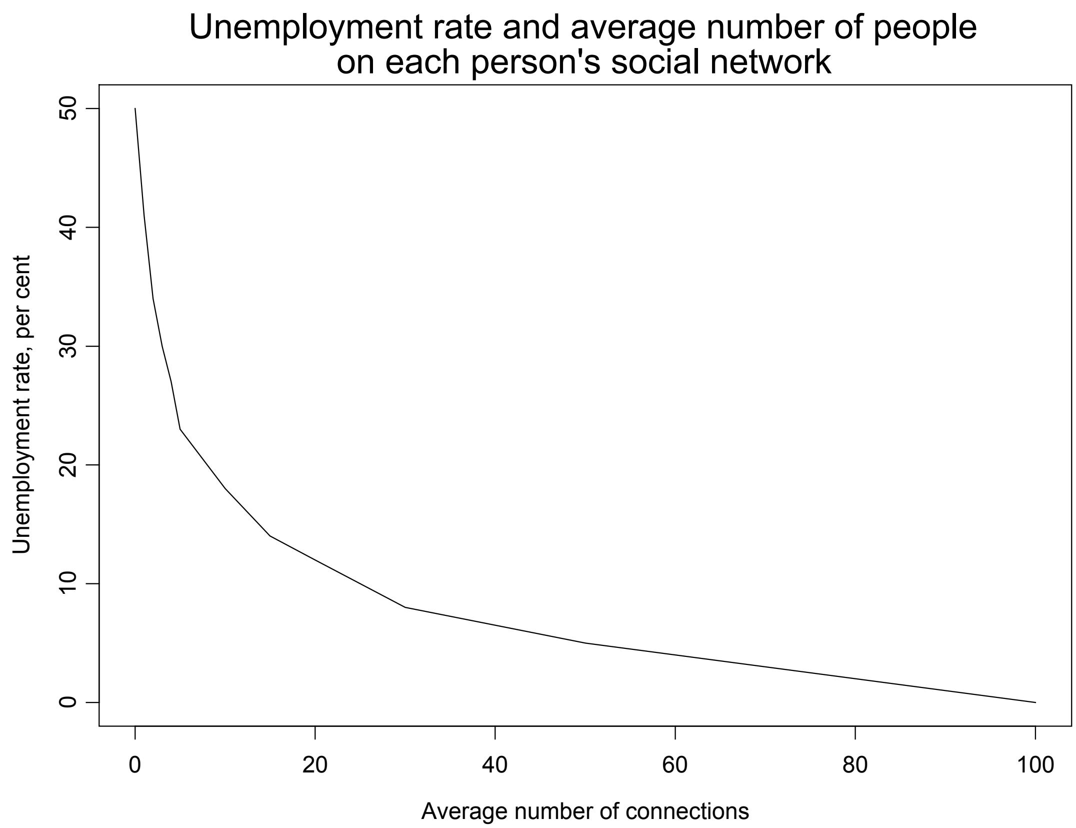

The interesting feature of the chart is that only a very small number of connections in a social network are needed to make a dramatic difference to the outcome. With each person being willing to pass on information to just five others, out of a possible one hundred, the unemployment rate falls by just over a half. This suggests that policies which succeed in increasing the flow of information about the availability of jobs in socially isolated communities can have a strong impact upon the rate of unemployment (provided, of course, that jobs are actually available). But in more integrated social groups, the return to such polices is much lower.

[6] The particular topology used to specify the network is a k-nearest neighbours one, where the agents are placed at random in 2-dimensional Euclidean space and each is connected to its k nearest

4. Conclusion

For most of the social sciences, the idea that individuals operate within society, and that their behaviour is affected directly by the way in which people are integrated, is a fundamental aspect of the discipline. In contrast, in standard economics, people operate in isolation, as autonomous individuals. The behaviour of any one person only affects others indirectly, in as much as it might change the set of relative prices which people face. Economics has been able to offer many insights into how the world works. But in order to do so, simplifying assumptions had to be made which precluded direct interaction between individuals. Many modern economists appear to have lost sight of the fact that the great, early formalisers of economic theory knew perfectly well that they were making assumptions which were dramatic simplifications of reality. But given the tools they had available, this had to be done.

Modern computing power enables more realistic assumptions to be made about how people behave, and the consequences of these assumptions analysed. The concept of formal modelling is not to everyone's taste, but the two examples above illustrate the potential of what can be done.

In the first, most of the assumptions of the ideal, perfectly competitive world of economic theory are retained. But simply by allowing firms to differ in their efficiency levels and then increasing the flow of information between them, we show that the prevailing market structure which emerges will be one of oligopoly, with markets dominated by a few large firms. This is the norm and not the deviation.

In the second, we offer an account of how high unemployment rates can persist in local areas even when the overall demand for labour is high. In disjointed communities, where the flow of information between individuals is low, unemployment can remain high even when jobs are being created. People simply do not get to hear about them in time.

Ultimately, it is only by making use of the insights of the other social sciences, and incorporating them in formal models such as these, that the economics community will be persuaded to move away from what has become a narrow and scientifically restrictive perspective.