Job Search, Unemployment and the Topology of Social Networks

Paul Ormerod and Laurence Smith[1]

Volterra Consulting Ltd[2].

November 1999

1. Introduction

The concept of job search features in both the economic and the sociological literature. An important aspect of this in economics is the existence or otherwise of imperfections in the transmission of information. Sociologists, whilst not necessarily sharing the economist's concepts of maximising behaviour and the desirability of perfect markets, have arrived at complementary kinds of ideas about the importance of efficient transmission of information.

Granovetter (1973, 1974) carried out much of the seminal sociological research on this issue. The key difference between the economic and sociological accounts is that the former stresses the characteristics of the individual, such as education levels, whilst the latter lays more emphasis on the social network in which individuals are embedded.

This paper sets out a way of formalising analytically the role which social networks play in job search and the persistence or otherwise of unemployment arising from information failure.

There is a great deal of empirical evidence to support the idea that exit from unemployment is to a large extent dependent upon social networks. Hannan (1999) cites, for example, Granovetter's original 1974 finding that over 60 per cent of professional and managerial workers in the US obtained their jobs through personal contact. Montgomery

[1] We are grateful to participants in a seminar organised by the Department of Sociology, University of Nottingham, and in particular to Robert Dingwall, for stimulating our interest in this issue.
[2] The Old Power Station, 121 Mortlake High Street London SW14 8SN; tel (+44) 0208 878 6333; pormerod@volterra.co.uk

(1992) reports many studies which essentially confirm this finding more generally in the US labour market. Hannan's own paper reports a detailed econometric study of micro-data held in the British Household Panel study, examining the impact of both traditional economic explanations of exit from unemployment, and the social network hypothesis. This latter is found to be very important.

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, or even the single most important, reason why unemployment exists. 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.

The overall approach could be applied much more widely to issues of social exclusion and the importance of the social network structure.

2. The model

The model 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 have a model populated by n individual agents. We might think of it as representing a housing estate or other local area. At any point in time, agents can be in one of two states of the world: unemployed (state 0) or employed (state 1).

All agents are identical. 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. If a job vacancy arises and an agent hears about it, this agent will apply for the job and be hired with certainty.

The model evolves through time in a sequence of steps. In each step, the following takes place. First, an agent is selected at random. If this agent is unemployed, he/she stays unemployed. If the agent is employed, the job is deemed to be lost, and the agent becomes unemployed. In other words, job losses occur at random.

This agent is returned into the general population, and another random choice of agent is made (in principle, this can be the same agent chosen in the first place, but unless n is small, the probability of this event is very low). A vacancy is assumed to have been created, which lasts for only a single period, and a single individual informed about it. If this agent is unemployed, he/she takes up the job and becomes employed.

If the selected agent is employed, the agent informs all the other agents which form part of his/her social network of the vacancy. If all these are employed, this step of the model ends. The vacancy is withdrawn, and is not carried over to the next period. If any of the agents on the social network are unemployed, one of them takes up the vacancy. The step of the model then ends.

Of course, these assumptions are not realistic. Agents may not apply for a whole variety of reasons. Even if they do so, the employer may not hire them. It is straightforward take account of such factors in the model by introducing a facility which allows different probabilities to be assigned to applying and to being accepted, both of which can vary between zero and one. The model can be further extended by allowing agents to differ in these probabilities. However, the purpose of this paper is to focus upon the role of social networks, and the assumptions are deliberately kept simple in order to do so.

The model can be summarised as follows. In each step, a random draw is made from the n agents, choosing agent i. If agent i is in state 1, with probability 1 he moves to state 0. If agent i is in state 0, he remains in state 0. A second random draw from all n agents is then made. If the agent chosen, agent j, is in state 0, with probability 1 he moves to state 1. If agent j is in state 1, all the other agents connected to agent j are selected. If they are all in state 1, they remain in state 1. If one of these agents is in state 0, with probability 1 the agent moves to state 1 (if more than one is in state 0, a separate random draw is made from this group to determine which one moves to state 1).

The question we explore in the paper is as follows. Suppose all agents start as employed, in other words in state 1. A solution of the model is obtained over m steps, where m is large. What are the implications for the rate of unemployment which emerges of a range of different topologies defining the connections between agents?

The model is very abstract. But it is not a completely unrealistic view of how the process of job loss and job creation might affect, say, the population of a single housing estate or tower block. The processes of job loss and creation impact on the area in ways over which the residents themselves have effectively no control.

3. Types of social network

In the first instance, we can distinguish the two extremes in terms of the degree to which agents are linked together on the social network.

First, suppose that each agent is connected to every other agent. An illustration of this, for only 6 agents to avoid the graphics becoming too dense, is set out in Figure 1.

Fully connected network with 6 agents

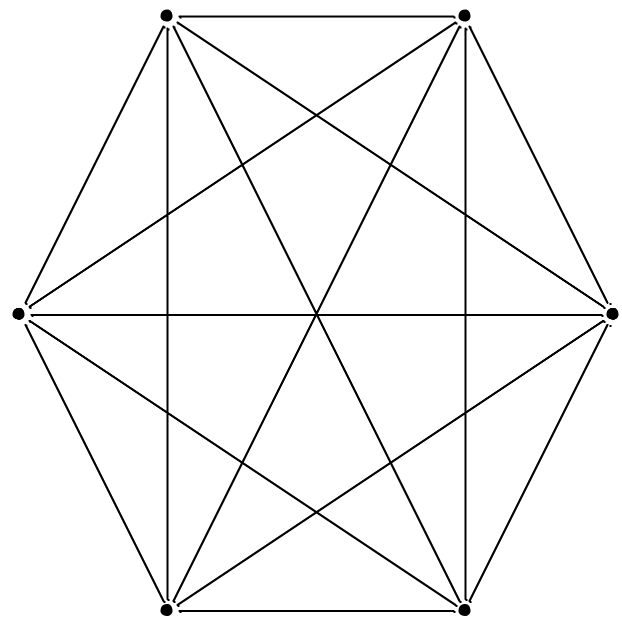

Figure 1

The solid black dots each represent one of the six agents, and the lines indicate connections between any agent, i, and all the other agents in the model.

In this form of network, information gained by any single agent is available by definition to all other agents. So all agents remain employed. In every step of the model, in the first instance an agent is chosen to lose his/her job. An agent is then chosen to be told about a job vacancy. If this agent is the one previously selected to become unemployed, he/she becomes employed again. If this agent is employed, all other agents, including by definition the one just made unemployed, will be informed of the vacancy.

At the other extreme, we can consider the case in which *no* agent is connected to any other. It is straightforward but slightly tedious (see Appendix) to show that as the number of agents increases, with this network the population as a whole approaches a 50/50 division between employment and unemployment.

The results obtained using other forms of network to define the connections between agents need to be placed in this context. In our artificial society which starts off with everyone in a job in a fully connected network, everyone is always in a job. And in a network with no social connections at all, unemployment averages 50 per cent of the total number of agents.

There is a wide variety of topologies which could be specified. There are three which have fairly straightforward interpretations. First, a torus. In this type of network, agents are placed at regular intervals on a grid. Figure 2 plots such a grid.

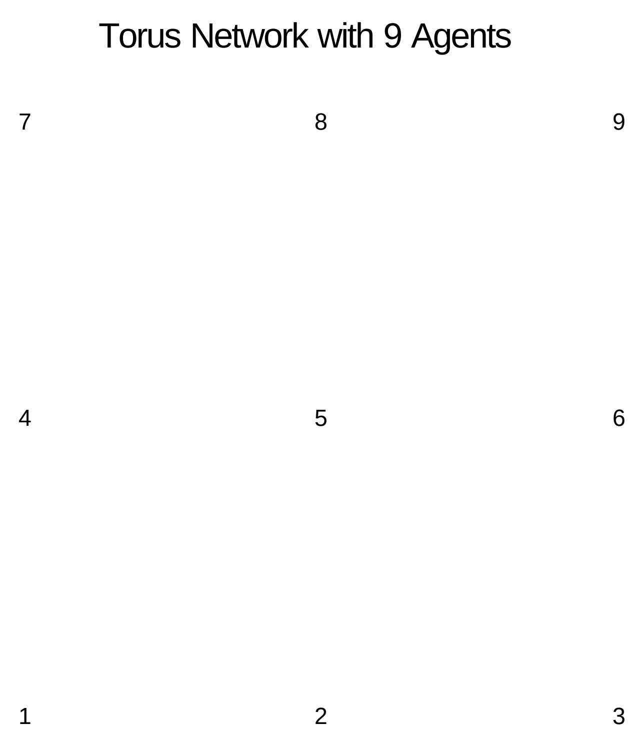

Figure 2

In this chart, the agents are labelled 1,2, etc. Each agent is connected to the four nearest other agents. So agent 5 is connected to agent 4 to the left, agent 8 above, agent 6 to the right, and agent 2 below. A further characteristic of a torus is that the grid can be thought of as being folded round itself, so that the edges connect with each other. As an example, agent 4 is connected to agents 1, 5 and 7. But it is also connected with agent 6, at the opposite end of the grid.

The concept of 'nearest' most obviously implies physical nearness, so that we might imagine that this kind of network represents agents learning about job opportunities from their next door neighbours. But the agents could be set out on the grid according to friendship or kinship structures rather than literally geographical nearness. This gives perhaps a more realistic interpretation of people learning from their immediate friends or family.

The second type of network, the k-nearest neighbour, is similar to the torus in terms of its interpretation, although there are some important differences. The agents are placed in two-dimensional Euclidean space at random, and each agent is connected to the k agents which are the shortest distance from him/her. The most obvious difference with a torus is that in the latter the number of connections per agent is fixed at 4, whereas with the k-nearest neighbour the number of connections, k, can take any value. As k is reduced towards zero, the k-nearest neighbours approaches the network with zero connections, and is k is increased towards n, it becomes more and more like the fully connected network

There is a second difference between the torus and k-nearest neighbour, which makes the latter somewhat more realistic in this context. On a torus, if agent i is connected to agent j, agent j is also connected to agent i. In other words, if agent i passes information to agent j, agent j does likewise to agent i. This symmetry need not hold for all agents in a k-nearest neighbour network. In general, for most agents this will be the case, but there will be a minority of cases in which this does not hold. In other words, agent j is a nearest neighbour of agent i and is therefore passed information by agent i, but agent i is not a nearest neighbour of agent j. Figure 3 illustrates this point.

K Nearest Neighbour Network, 10 agents, 2 connections

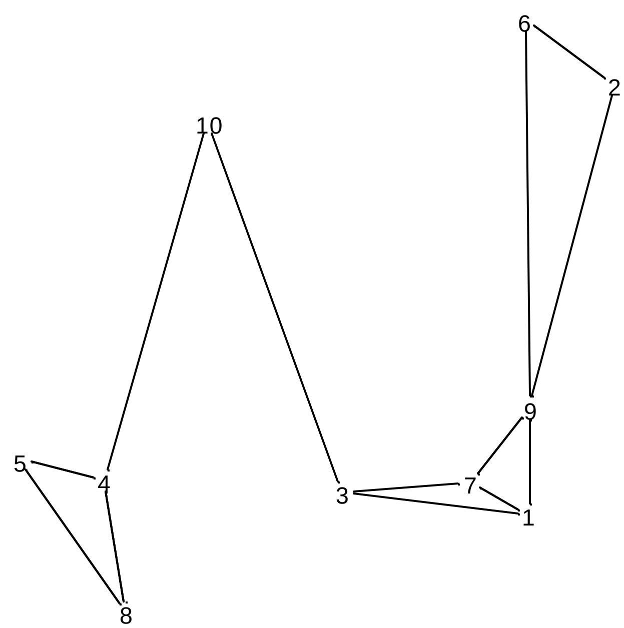

Figure 3

The two nearest neighbours of agent 9 are clearly agents 1 and 7, and of agent 6 they are 2 and 9. So agent 6 passes information to agent 9, but agent 9 does not do so to agent 6.

The final network we examine is that of a tree. Agents differ in the number of other agents to which they are connected. But with a tree, the majority of agents are usually only ever connected to a single other agent. They are, as it were, stuck out on the end of a branch. Figure 4, with an illustrative tree network, makes this clear

Illustrative tree network

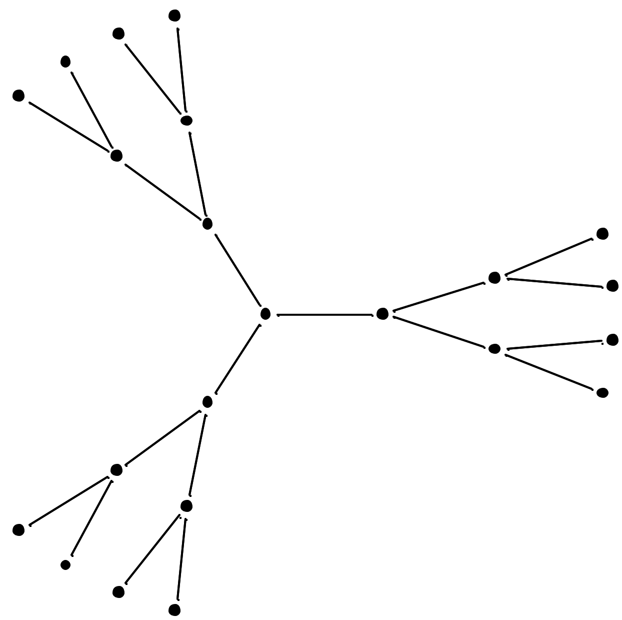

Figure 4

.

The tree geometry can be thought of as an area in which agents are not particularly well connected to each other socially. The maximum number of connections per agent is small relative to the total number of agents. Further, the majority usually have very poor connections - being linked to only one other agent[3].

4. The results

Before reporting more detailed results, two general points need to be made. First, the results are not sensitive to the initial conditions of the agents' states of the world. We postulate that all agents begin in state 1, in other words in employment. But, providing that the models are solved over a sufficient number of steps, the eventual outcome is the same even if they all start off in state 0. Second, the results scale over a reasonable range of values for the total number of agents in the model. In other words, the proportions of the total population in each state which emerge are not in general sensitive to the total number of agents, n. For very small values of n (around 10 or less), this is not the case, but otherwise this result obtains for n up to around 100,000 (larger numbers were not investigated because of the length of time required for solutions).

In each of the cases reported here, 100 separate solutions of each particular specification of the model were run over 2,000 steps. The average number of agents in each state of the world was calculated in each solution for steps 1,500 to 2,000.

In section 3 above we noted two analytical results which can be readily obtained for the two extreme cases of the degree of connectedness. In a network where each agent is connected to every other agent, all agents are employed. In a network in which no agent is connected to any other agent, the proportions of employed and unemployed approach a 50/50 division as the number of agents increases (even for as few as 20 agents, the split is 51.3/48.7, and for 50 it is 50.5/49.5).

The results need to be placed in this context. In other words, unemployment in this model can vary between zero and 50 per cent, depending upon the nature of the connections between agents.

With both the torus and tree networks, the number of connections per agent relative to the total number of agents is in general very small. In consequence, high rates of unemployment emerge. Information about job vacancies is passed to very few other agents.

---

[3] Let ρ be the number of branches and d the number of connections for each agent not on the outside edge of a branch, then the number of agents on the outside is $(d - 1)^{\rho - 1}.d$, and the total number of agents is 1 plus

When the social network is specified in terms of a torus, the model settles down around an average rate of unemployment of 25 per cent. This is the average over 100 separate solutions of the model, as mentioned above. Because of the probabalistic nature of the model, each solution of the model is unique. Over the 100 solutions as a whole, the lowest average rate of unemployment which emerges is 16 per cent, and the highest is 35 per cent.

With a tree network, in which the average number of connections per agent is less than in a torus, the average rate of unemployment is higher. It is slightly sensitive to the exact specification of the tree, but over a range of different tree structure, the rate of unemployment averages between 32 and 36 per cent - although just as with the torus, there is a range around any given average of the outcome of a large number of individual solutions.

An illustrative outcome of the model with a tree network is given in Figure 5. For clarity of the graphics, the number of agents is kept fairly small (46), and they are indicated by the state of the world in which they are in after 1000 steps of the model. The symbol "e" denotes an agent in employment, and "u" and agent who is unemployed.

---

the sum over i from 1 to ρ of $(d - 1)^{i-1}.d$

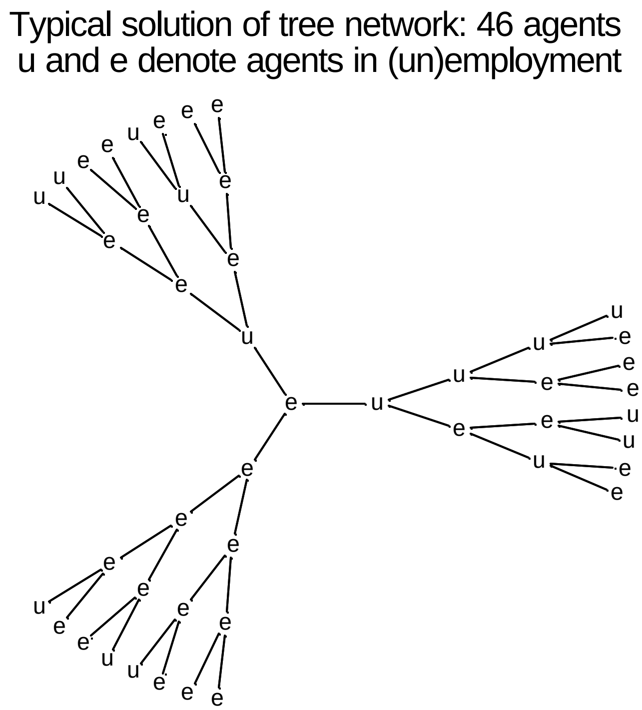

Figure 5

In this snapshot of a particular step in the model solution, 15 out of the 46 agents are unemployed, or almost 33 per cent.

The k-nearest neighbour network can be investigated in more detail. Table 1 shows the average rate of unemployment over 100 solutions of the model (averaged from steps 1,500 to 2,000) for this type of network with different numbers of connections per agent. There are 100 agents. It must be remembered that in this network, the connections between agents are not always symmetrical ie: agent i may pass information to agent j, but agent j need not pass information to agent i. The results for k=4, for example, therefore differ slightly from that of the torus, where k is also equal to 4 but where connections are symmetrical.

Table 1. Average rate of unemployment, 100 solutions of the model

k-nearest neighbour network

k	rate of unemployment
0	50
1	41
2	34
3	30
4	27
5	25
10	18
15	14
20	12
30	8
50	5
100	0

Compared to the maximum possible rate of unemployment in this model, the rate falls rapidly once even a small number of connections between agents is introduced. 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.

4. Conclusion

This paper sets out a deliberately stylised model of the evolution of unemployment amongst small social groups, such as an individual housing estate or even a single tower block on such an estate.

The jobs of individual agents in the group are lost at random. New ones are also created at random, but information about the availability of these new jobs is limited. Individual agents are connected on a variety of different geometries, and pass information about job vacancies to other agents to which they have a direct connection.

In this model, unemployment can persist only because of a failure of the social network to inform individuals about job opportunities. Of course, in reality this is far from being the case. But where such a problem does exist, the model highlights the importance of the degree of connectedness, as it were, of the social network.

High rates of unemployment can emerge in social networks which are only weakly connected, and the returns to policies which can increase the degree of connection are likely to be high.

The approach can be used more generally to investigate the importance to social exclusion issues of the particular ways in which social networks exist.

Appendix: Evolution of employment in a discrete network

Let x_t be the number of agents employed at time t, and n the total number of agents. No agents are connected.

Consider the process which generates x_{t+1}:

A single agent is chosen at random in two separate draws in each time period. The possible outcomes are as follows:

i) the agent in the first draw is employed and the agent in the second is employed

ii) the agent in the first draw is employed and the agent in the second is unemployed

iii) the agent in the first draw is unemployed and the agent in the second is employed

iv) the agent in the first draw is unemployed and the agent in the second is unemployed

In (i), $x_{t+1} = x_t - 1$; in (ii) and (iii), $x_{t+1} = x_t$; in (iv), $x_{t+1} = x_t + 1$

Multiplying these outcomes by their respective probabilities, we obtain:

$x_{t+1} = x_t$ [$(x_t - 1)((x_t^2 - x_t)/n^2) + (x_t^2/n^2)(2n - 2x_t + 1) + (x_t + 1)((n - x_t)^2/n^2)$]

Letting $t \to \infty$, if a limit exists, $x_{t+1} = x_t$, which implies that

$x_t = n^2/(2n- 1)$, so that the proportion of agents in employment is $n/(2n - 1)$

References

M.Granovetter, 1973, 'The Strength of Weak Ties', *American Journal of Sociology*, vol. 78, pp.1360-80

M.Granovetter, 1974, *Getting a Job: A Study of Contacts and Careers*, Harvard University Press

C.Hannan, 1999, 'Beyond Networks: 'Social Cohesion' and Unemployment Exit Rates', University of Essex Institute for Labour Research, Discussion Paper 99/28

J.Montgomery, 1992, 'Job Search and Network Composition: Implications of the Strength-of-Weak Ties Hypothesis', *American Sociological Review*, vol.57, pp586-96