Social Networks and Access to Financial Services in the UK

Paul Ormerod (pormerod@volterra.co.uk) *

and

Laurence Smith (lsmith@volterra.co.uk)

Volterra Consulting Ltd

5 The Old Power Station

121 Mortlake High Street

London SW14 8SN

March 2001

The work for this paper has been financed by the Britannia Building Society, as part of the project on Social Exclusion from Financial Services

* Corresponding Author

Abstract

Empirical evidence for the UK shows that almost all individuals without access to financial service accounts are not in work. But amongst the total population of those not in work, the economic and social characteristics, such as age, income, gender and geographical location, of those with and without financial service accounts are very similar.

We set up a theoretical model based on graph theory to see to what extent it is possible that the different social networks of individuals can contribute to whether or not they have accounts. Using abstract but nevertheless realistic characterisations of networks, we find that the composition of the group from which an individual is willing to take advice on financial services can exercise a strong influence on behaviour. Individual agents without accounts tend to have a much higher proportion of individuals also without accounts on their particular social networks than does the population as a whole. This result is supported strongly by empirical evidence

1. Introduction

This paper considers the application of graph theory to the issue of individual agents who do not have access to financial services. In other words, these are agents who do not have an account with any financial institution. Using the current policy jargon in the UK and Europe, we examine the question of social exclusion from financial services.

The motivation for the paper arises from the characteristics of those individuals without financial service accounts ('the excluded') described in Meadows and Ormerod [1]. Information is provided on around 42,000 individuals in the 1997/98 Family Resources Survey and 3,400 individuals in the Office of National Statistics Survey in March and April 2000. A characteristic shared by virtually every single excluded individual is that he or she is not in work. This covers a wide range of circumstances, such as the unemployed, the retired, and those on invalidity benefit.

Standard logistic regression is used on the whole sample of all individuals not in work for both surveys separately, to examine the extent to which the possession or otherwise of a financial services account can be explained by differences in economic and social characteristics of these individuals, such as age, income, gender and geographical location. In common with previous studies of this question (for example, [2,3,4]), Meadows and Ormerod find that logistic regression models have only weak levels of explanatory power. In other words, other factors are required in addition to the social and economic characteristics of the individuals in order to account for financial exclusion.

This paper offers a theoretical model which analyses exclusion from financial services on the basis of the social networks, or graphs, which connect individual agents. It is postulated that the decision of an agent whether or not to start a financial account is influenced by the recommendations of other agents whose opinion the agent values on this matter.

Such agents can be thought of as being connected on a graph. A graph is a collection of points, or vertices, with lines, or edges, connecting pairs of them. Empirical analysis of other types of social network suggest quite clearly that their properties require the individual agents in the network to be connected in non-random ways. Wilmott et.al. [5] and Foster et. al. [6], for example, report findings on friendship networks, Watts and Strogatz [7] the connections between movie actors, and Davis and Greve [8] the interlocking structure of membership of corporate boards.

Section 2 of the paper sets out the basic theoretical model. Section 3 describes a range of different topologies by which agents might be connected. Section 4 describes the results from running the model on these different topologies, detailing the effect on the percolation of financial services amongst individual agents. Section 5 sets out the conclusion.

2. The theoretical model

The empirical evidence suggests that almost all agents without access to financial services are not in work. But the economic and social characteristics of all individuals not in work provide only a low amount of power in discriminating between those with and without accounts. It is therefore not unreasonable in this context to assume that the n individual agents who populate the model have identical social and economic characteristics.

The empirical importance of social networks in determining whether or not an individual not in work does or does not have an account is summarised in Table 1.

Table 1: USE OF FINANCIAL SERVICES BY FAMILY AND FRIENDS (non-working sample)
column percentages

Use of accounts by friends and family	Proportion of people who have accounts	Proportion of people who do not have accounts
All or most have accounts	87	38
Some have accounts	6	26
Few or none have accounts	2	14
Don't know	6	21
TOTAL	100	100

Source: ONS Omnibus Survey March/April 2000

The null hypothesis that the use of accounts by family and friends is the same between those who do and do not have accounts is rejected very decisively.

In the model described here, the only factor which determines whether an individual will decide to set up an account is a recommendation from another agent whose opinion he or she respects on this matter. In other words, we abstract from differences between agents' social and economic characteristics. Further, we abstract from the supply side of financial services, and assume that an agent who applies for an account is automatically given one. This is not completely realistic, but the purpose of this paper is to see how patterns of financial exclusion can arise purely on the basis of the flows of information (recommendations) between agents.

The agents in the model are connected to each other by a network, and differ only in terms of the composition of the network. An agent is connected only to those agents whose opinion is taken into account in deciding whether to have a financial account.

In terms of the formal model, an agent can be in one of two states of the world. He or she can either not have an account (state 0), or have an account (state 1). To begin with, all agents are in state 0, and an individual is chosen at random to move to state 1.

The model then evolves in a series of discrete steps. In each step, an agent is drawn at random to consider whether or not to change his or her state of the world. If the agent is in state 1, he or she remains in state 1, and the model moves onto the next step. An implication of this is that agents never move from state 1 back to state 0. Whilst this is not a completely accurate description of reality, it is nevertheless a very good approximation to it. Once someone has an account, it is rare to move back to not having one.

If the agent is in state 0, the agent examines the state of the k individuals whose opinions he or she respects in this matter. At any particular step of the solution, m of these will be in state 1, where $0 \leq m \leq k$. The agent decides to move immediately to state 1 with probability m/k. The model then proceeds to the next step, where another agent is drawn at random to decide. The draw is done with replacement.

The model evolves step by step until 90 per cent of all n agents in the model are in state 1. Around 90 per cent of all individuals out of work in the UK have financial accounts [1], in other words are in state 1 of the world. Due to the randomness in the generation of certain types of networks which connect individuals, it may not be possible for 90 per cent of the agents to end up in state 1. Solutions where this is the case are ignored.

For each kind of network, a large number, S, of separate solutions of the model are obtained, and the properties of these S simulations examined.

3. Network Topologies

Our motivation in the paper is to explore the possibility that the social networks of individuals can be important in determining whether or not agents have financial service accounts. The phrase 'social network' in this context does not, of course, mean the general network of family, friends and acquaintances. It means the specific group of people whose opinion an individual takes into account in deciding whether or not to take up a financial account.

We set out first of all three examples of networks which can be thought of as benchmarks with which to compare others, the discrete, the completely connected and the random network. The discrete and completely connected networks are in fact special cases of the more general random network, and can be described briefly.

3.1 Discrete Network

Agents can be connected by a wide range of different networks. At one extreme, we might imagine a network with no connections at all. In this world, no agent pays attention to the opinions of any other, and there are no interactions.

3.2 Complete (or Fully Connected) Network

At the other extreme, agents might be connected by a network which links each agent to every other agent. In the first step of any particular solution of the model, one of the n agents is chosen at random to move from state 0 to state 1. So in the second step, providing this same agent is not the one chosen again, the agent chosen considers the state of the world of all other agents and observes that just 1 out of the other n-1 agents has an account. This agent will therefore take up an account with probability 1/(n - 1).

It is obvious that with this kind of geometry connecting the agents, all agents will eventually end up in state 1.

3.3 Random Network

Of more general interest is the case of a random network. In this case, each of the agents is connected to any other agent with a fixed probability, p. If agent i is connected to agent j, this means that agent i takes account of j's opinion in deciding whether to move from state 0 to state 1.

In a random network, if agent i is connected to agent j, agent j will not necessarily be connected to agent i. Indeed, j will be connected to i with probability p. This seems to be a more realistic property of a network in this context than the assumption that if i pays attention to j, j will necessarily listen to i.

It is easy to see that as p approaches zero, the random network becomes increasingly similar to the discrete network, in which no connections exist at all. As p approaches the value of one, the random network becomes like the completely connected network.

A formal way of expressing the difference between networks is by the clustering coefficient [7]. The clustering coefficient for an agent is the fraction of total possible connections between its neighbours that actually exist. In other words, the clustering coefficient measures the degree of overlap between the sets to which agents are connected. For example, if agent a is only connected to b and c, then the total number of possible connections between a's neighbours is 2 (b connected to c and c connected to b). If b is connected to c but c is not connected to b then a's clustering coefficient is 0.5.

The clustering coefficient for the network, C, is the average of all the agents' clustering coefficients. So for the completely unconnected network, C = 0. For the completely connected network it is 1, and for the more general random network it is equal to p.

3.4 Small World Network

As noted in the introduction, random networks have been studied intensively, and they have empirical application. However, in many social contexts, they lack an important aspect of reality. The probability of two agents being on the same social network is usually higher if they are both connected to the same third agent, than if they are not. If agents i and j are both connected to agent k, the probability of i and j being connected is higher than if they were not connected to k. The clustering coefficient of a random network is p, the probability of a connection existing from any given agent to any other given agent, but in many social contexts the clustering coefficient is greater than p.

Watts and Strogatz [7] introduced the concept of a small world network in 1998. . An illustrative small world network is shown below. There are a number of variants of the small world concept, but Figure 1 illustrates the main features.

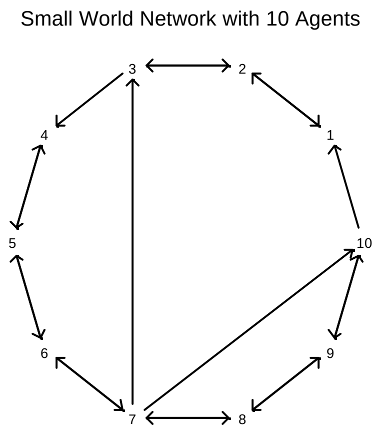

Figure 1

In a small world network, the agents are initially arranged in a ring with every agent being connected to k agents on each side of it (in the figure k=1)[1]. The agents either side of any given agent can be thought of as his or her immediate neighbours in a social context, or the agents to which he or she feels closest.

Figure 1 is plotted with k = 1 for ease of graphical exposition. However, in the case where k = 1, the clustering coefficient is zero. Agent 5 in Figure 1, for example, is connected to agents 4 and 6. In theory, agent 4 could also be connected to agent 6, and vice versa. But with k = 1, none of these links exist, and hence the clustering coefficient is zero. It is apparent, however, that if $k > 1$, the clustering coefficient becomes decidedly non-zero. If k = 2, for example, agent 3 would be connected to agents 1, 2 , 4 and 5. Agent 2 would be connected to agents 10, 1, 3 and 4. The clustering coefficient is 0.5. In general, each agent is connected to 2 k other agents and the clustering coefficient is equal to $3 * (k - 1) / (2 * (2 k - 1))$, (provided the total number of agents is greater than 2k) so that C increases as k increases.

Figure 1 illustrates an example of an important extension of the small world concept. The agents are placed on the ring, and with a given probability p, each connection is randomly rewired. In other words, the connection from an agent to its immediate neighbour is removed, and a random one replaces it. So, for example, in figure 1, the connection from agent 1 to agent 10 is removed, and at random a connection from agent 7 is inserted to agent 10. Agent 10 therefore pays attention not to agents 1 and 9, its immediate neighbours, but to agents 7 and 9.

An alternative extension is to make further connections across the ring without removing the immediate connections. Agent i will be connected to agent m, which is not one of its k-nearest neighbours, with probability p. In this context, the analytical interest is in how many further connections need to be added for the network to approximate that of a completely connected network - a 'small world' in which everyone is connected to everyone else.

Small world networks are a fascinating and important set of networks. However, a potential drawback of the concept is that for $k > 1$, the clustering coefficient of the initial network is high. For some real world networks, this might be desirable. Defining Hollywood movie actors as being connected if they have appeared in the same film a

[1] Moore and Newman [9] examine networks in which not all agents on the ring are connected initially.

clustering coefficient of 0.78 is found [7]. A study of friends, relatives and neighbours in the UK [6], however, found a clustering coefficient of only 0.34, and cited several other similar kinds of study in which coefficients of between 0.16 and 0.44 were reported.

The clustering coefficient can be reduced for small world networks once random connections are introduced, with the number of local connections being reduced pro rata. By way of illustration, consider a small world network of this kind populated by 100 agents, each connected to the two nearest neighbours on each side (k = 2 in the context of the discussion above). The value of C is 0.5. Setting the probability, p, to 10 per cent that agent i is not connected to agent j, where these are near neighbours, but is instead connected to another agent at random, results from simulations of the network give a clustering coefficient which averages around 0.37.

4. Results of the theoretical model

Our interest in the model is to see to what extent a social network might influence the financial exclusion of agents. In any particular simulation of the model, each agent starts off in state 0 (ie: without an account), and an agent is chosen at random to switch over to state 1. The percolation of agents into state 1 is monitored, and the simulation ended when either the percentage of all agents in state 1 is 90, or the model has gone through 1000n solution steps, where n is the total number of agents. A 90 per cent penetration of account holders was chosen because it represents reality [1].

In those solutions in which 90 per cent reach state 1, we examine the state of the world of the agents who are connected to those agents who remain in state 0. By carrying out a sufficient number of independent solutions of the model for any given social network, we obtain an estimate of the distribution of this outcome. We ignore solutions which do not reach the 90 per cent threshold for agents in state 1.

4.1 Discrete and Complete Networks

The outcome in both cases is obvious. In the discrete network, since there are no connections between agents, no agents have neighbours with accounts and so the number of agents in state 1 never increases above the one agent drawn at random in the first step.

For the completely connected network, since all agents are by definition connected to all others, the answer is the same for all agents. Thus, for each agent in state 0, the proportion of all agents on his or her network which is also in state 0 is 10 per cent, and the proportion in state 1 is 90 per cent.

This implies that if the social network is of this latter kind, it can play no role in explaining whether or not agents remain in state 0 once we place it in a more general setting along with the social and economic characteristics of the agents. The network of each agent in state 0 is identical for all such agents.

More importantly, if the relevant networks approximate the properties of a completely connected network rather than being absolutely identical, a network effect is unlikely to offer much explanatory power. Such an outcome would arise, for example, with a random network with a high level of p or with a small world networkwhere the total number of connections is very high.

4.2 Random Network

In the case of the discrete and completely connected networks, the results can be obtained by simple logic. With other kinds of networks, however, analytical results are either very hard or impossible to obtain. We therefore relied upon computer simulations of the model.

The precise outcome will of course depend upon the particular specification of the network. In this case we set the probability of agent i being connected to agent j at 0.05.

In other words, any given agent is connected to, on average, 5 per cent of all agents in the model.

A small number of simulations were carried out with 10,000 agents and compared to those obtained with just 100 agents. The results scale linearly, so for ease of computation in these and subsequent results reported we used 100 agents.

Once 90 per cent of agents are in state 1 in any particular simulation, we can then consider those that remain in state 0. In general we find that the agents to whom those in state 0 are connected have the following properties. First, the proportion of agents to which a state 0 agent is connected who are in state 1 is less than 90 per cent. And, second, the proportion of those agents in state 0 is in general more than 10 per cent. In other words, there is a certain amount of potential explanatory power of the social network in this case. Those who remain in state 0 tend to have more neighbours in state 0 than is implied by the population average, and fewer in state 1.

We performed 500 simulations of random networks with 100 agents, with an average of 5 connections per agent. We calibrated the percentage of a non-account holding agent's acquaintances (ie: those to whom the agent is connected) who have accounts, so that "few" is if less than or equal to 25% of the acquaintances have an account; "some" is if greater than 25% and less than 75%; and "many" is if greater than or equal to 75% of the acquaintances have an account. The results are:

- 5% of agents without accounts have few acquaintances with accounts
- 21% " " have some acquaintances with accounts
- 74% " " have many acquaintances with accounts

By contrast, in a completely connected network, *all* agents without accounts have 90 per cent of acquaintances with accounts and 10 per cent without. In other words, all agents without accounts have 'many' acquaintances with accounts. In the random network in this example, only 74 per cent fall into the same category.

This difference in the properties of a completely connected and a random network is not dramatic, but it is a very important one conceptually. It illustrates the possibility that a key reason why certain individuals do not take up financial services is because of their social networks or, rather, the networks of individuals from whom they are willing to take advice regarding financial services.

This property emerges much more clearly in a more realistic geometries, namely that of a small world network.

4.3 Small World Network

We explored a number of different specifications of this topology, and report here the results for one which generates an outcome similar to the empirical evidence set out in Table 1.

Each agent is initially connected to its four nearest neighbours (two on each side round the ring), with a probability p = 0.2 of rewiring, we have a network with a similar clustering coefficient to [8] and each agent has 4 acquaintances. Performing 500 simulations of the model and examining the acquaintances of those without accounts we find:

- 57% of agents without accounts have few acquaintances with accounts
- 26% of agents without accounts have some acquaintances with accounts
- 17% of agents without accounts have many acquaintances with accounts

We can see that the results are different to a random network. In the random network, the distribution of numbers of acquaintances (of people without accounts) with accounts was skewed so that the people without accounts had a high proportion of acquaintances with accounts. For the small world network, there is a much higher proportion of agents who have either few or some acquaintances with accounts, and correspondingly, a much lower proportion of agents who have many acquaintances with accounts. This suggests that the type of social network can have a large influence on the distribution of account holders.

In fact, comparing these numbers to Table 1 (after removing the "Don't know"s from the statistics), the numbers are similar. From Table 1, the percentages for most, some and few friends and family having accounts are 48%, 35% and 18% respectively. These are clearly similar to the 57%, 26% and 17% from our model.

.

5 Conclusions

Empirical studies of people with and without financial accounts find that a common characteristic of almost everyone without an account is that he or she is not in work. Apart from this, however, statistical models which attempt to distinguish between those with and without accounts on the basis of the social and economic characteristics of individuals have only a low degree of explanatory power.

In this paper, we set up a theoretical model which examines the potential impact of social networks on social exclusion from financial services. In other words, we investigate whether the social network of an individual can in principle affect the decision or otherwise to take up a financial services account.

The model is deliberately abstract, and in order to focus on the role of networks, the only reason in the model why an individual takes up a financial service account is on the basis of whether or not the people from whom the individual is willing to take advice on this matter themselves have accounts. The higher the proportion of such agents with an account, the higher the probability of any given agent without an account acquiring one at any point in time.

The set of people from whom any given individual will take advice will often be completely different for different people, and the networks which we examine reflect this fact.

Even in a network in which the connections between individuals are purely random, we find that the outcome permits an influence of the social network on the state of the world of any individual, that is whether or not he or she has an account. Those without accounts tend to have a higher proportion of agents on their networks also without accounts than is typical of the population as a whole.

This finding is much clearer when a more realistic type of geometry connecting agents is considered, namely a small world network.

In summary, the network of individuals from whom any given agent is willing to take advice on the decision whether or not to take up a financial services account can in principle have a powerful affect on the outcome and hence on the persistence of social exclusion from financial services.

References

[1] P.Meadows, P.Ormerod and W.Cook (2000), *The Role of Social Networks in the Consumption of Financial Services: Technical Report*, National Institute of Economic and Social Research, London

[2] E.Kempson and C.Whyley (1999), *Kept Out or Opted Out: Understanding and Combating Financial Exclusion*, Policy Press for Joseph Rowntree Foundation, York

[3] Office of Fair Trading (1999), *Vulnerable Consumers and Financial Services: Report of the Director General's Inquiry*

[4] E.Kempson and C.Whyley (1998), *Access to Current Accounts*, British Bankers Association, London

[5] P.Willmott, *Friendship Networks and Social Support*, Policy Studies Institute, London

[6] C.C.Foster, A.Rapoport and C.J.Orwant, (1963), 'A Study of a Large Sociogram: Elimination of Free Parameters', *Behavioural Science*, vol.8, pp.56-65

[7] D. J.Watts and S. H. Strogatz, 'Collective Dynamics of "Small-World" Networks', *Nature*, vol. 393, 4 June 1998

[8] G.F.Davis and H.R.Greve (1997), 'Corporate Elite Networks and Governance Changes in the 1980s', *American Journal of Sociology*, vol.103, pp.1-37

[9] C.Moore and M.E.J.Newman (2000), 'Epidemics and Percolation in Small-World Networks', cond-mat/99 11492 v2