Innovation, Diffusion and Agglomeration

Paul Ormerod[1], Volterra Consulting, and Bridget Rosewell[*2],
Greater London Authority Economics

Paper to be presented at the Brisbane club meeting, June 2007

……

pormerod@volterra.co.uk
Bridget.Rosewell@london.gov.uk
*   corresponding author

Abstract

Innovation takes place in time and space – this has been neglected in more traditional accounts.

New economic geography has drawn attention to the mechanisms which support cities in a static maximising framework

This neglects on the other hand growth and innovation. Yet we observe that innovation happens particularly in cities and that productivity, which is associated with innovation, is higher in cities.

What is it about cities which is important? Density, networks, diffusion. Dense networks can be created quite easily?

Diffusion across these much more effective.

1.   Introduction

Innovation, in all its variety of forms, is a necessary ingredient in growth and in improvements in productivity. For Adam Smith, it was the division of labour and continued improvement in that division which drove growth. Later the focus has moved to technological change, process and product development, and most recently knowledge and learning. These have been characterised by Antonelli[1], as deriving from major strands in economic thought from Smith through Schumpeter and Arrow to Marshall.

However, the characterisation of types of innovation and their incorporation into actually existing economic activity has also to locate that activity in time and space. This applies as much to the ability to create divisions of labour as it does to the more abstract benefits of knowledge spillovers.

Location in time and space in turn has implications for the potential for innovation to become embedded in business as usual or to create further innovative possibilities. This is the context in which the economics of city existence and development becomes important to any story of how innovation takes place.

---

[1] C Antonelli, (2007) The Foundations of the Economics of innovation: from the Classical Legacies to the Economics of Complexity', paper circulated to Brisbane Group

The role of cities has perhaps been understated or even forgotten in the technology and energy focus of how innovation has proceeded. The well worn story of the industrial revolution is based on harnessing water, coal and steam to produce energy. Cities grew up around these activities but were viewed as accidents of co-location. This is misconceived. The majority of the population on the planet now lives in cities, and growing economies are growing cities still faster. Innovation and cities appear to be a key part of the feedback mechanisms underlying economic growth and productivity performance.

Dosi(1997)[2] identifies four elements of analysis of innovation – opportunities, incentives, capabilities and organisational arrangements and mechanisms. However, these are thought of as applying largely to the firms who incorporate innovation, although other kinds of organisations such as universities are mentioned. The wider interaction between different kinds of firms and other bodies once they are co-located does not seem to appear in the literature.

2.   The Role of Cities

Every civilisation has produced cities and the collapse of civilisations has equally been marked by a shrinking of city life. The end of the Roman Empire saw cities shrink where they did not disappear, and stone buildings be replaced by wood and thatch. Public buildings fell into disuse as the administration, tax collection, long distance trade and military functions all disappeared. Although historians now chart a more complex process of economic difficulty and re-emergence than the blanket term 'Dark Ages' might suggest, and although different parts of Europe experienced different trajectories, nonetheless it was not until the eleventh century that those cities which had survived once again covered a similar land area to their Roman counterparts[3].

And as the cities recovered, so did economic performance. Long distance trade began to focus on and create larger cities, while regional and local trade created smaller centres.

---

[2] G Dosi, Opportunities, Incentives and the Collective Patterns of Technological Change, Economic Journal, 107, 1530-1547, 1997
[3] D Nicholas, The Growth of the Medieval City, Pearson, 1997

Finance and banking were developed alongside artisanal craft and luxury production. Early medieval Europe is the story of the revival of cities and of trade, rather than of major technological breakthroughs. The invention of horseshoes, the horse collar and stirrups played a role – if only in military conquest – but are now thought to have been of less importance to the revival of agriculture and population except in some limited circumstances.

The revival of cities in medieval Europe is one story. Many conditions have changed since then. In particular, we have seen the introduction of the capitalist system which has shown itself to be capable of supporting much larger numbers of people at a higher quality standard of living than any previous economic system. The importance of cities to the economy remains unchanged. In the USA, 75 per cent of people live in cities. There are now 20 cities in the world of more than 10 million people, compared to only 2 in 1950. The UN estimate that just over half the population lives in cities across the globe.

Recent estimates by PwC[4] suggest that Mexico City is the highest ranking developing city in output terms, and they expect Shanghai, Mumbai, Istanbul Beijing and Manila to enter the ranks of the top 30 cities over the years to 2020. The top 30 cities by GDP are currently estimated to account for around 16% of world GDP, rising to 25% for top 100 by 2020. The PwC projections show the pack continuing to be led by the existing world cities of New York, Tokyo, Los Angeles, London, Chicago and Paris.

The overwhelming conclusion is that under any economic system, cities are important and moreover cities' growth and success are associated with economic growth and improved performance. Since improved performance means innovation and its diffusion, cities must contribute to this. However, this aspect of understanding innovation has been largely neglected, perhaps because the subject of city growth and has been thought of as the purview of geographers rather than economists. By contrast, economists have focused more on understanding the nature and scope of enterprises and have therefore concentrated on the investigation of how innovation is created and diffused through firms and entrepreneurs. This paper argues that properly to understand

---

[4] PricewaterhouseCoopers, Economic Outlook, March 2007

the process of innovation, we need to focus on its location as well as its organisational base.

There has been a revival in recent years of interest in economic geography, a research agenda labelled the new economic geography has applied modern models of profit maximising firms to a world which includes transport costs for both goods and workers[5]. Such models show a rationale for the existence of cities in a static framework and indeed identify the scope for feedback mechanisms including knowledge spillovers and scope for learning. However, little of this literature explicitly incorporates innovation or growth into the models. Where they do, such as Duranton and Puga[6], the model explicitly looks for an equilibrium solution for profit maximisation, although there is explicit room for connections and experimentation.

This is a serious omission, since the ability to innovate and manage innovation is an important aspect of city performance. According to the US Small Business Administration Innovation Database, 45 per cent of the innovations in 1982 were concentrated in four cities: New York, San Francisco, Boston and Los Angeles. Fewer than 4 per cent of the innovations occurred outside metropolitan areas[7]. Although this data is now quite old, it is complied from trade journals and other detailed sources to provide a picture of new products in manufacturing. It seems likely that innovation in services might be still more concentrated than that in manufacturing. Feldman and Audretsch conclude from this data than innovation was spurred by diversity and spillovers between industries. Again, this is still more likely in business and financial services which are heavily centred in cities.

The focus of the study of innovation has on the other hand concentrated on enterprises and on industries. This is not very surprising given that firms are where innovations have to be realised. But at the same time this has meant neglecting the role of location,

---

[5] See for example, M Fujita, P Krugman and A J Venables, "The Spatial Economy", MIT, 1999; M Fujita and J-F Thisse, Economics of Agglomeration, Cambridge, 2002

[6] G Duranton, D Puga, 'Nursery Cities,: Urban Diveristy, process innovation and the life cycle of products, American Economic review, December 2001, p 1454-1477

[7] M P Feldman, D B Audretsch, (1999) Innovation in Cities: Science based diversity, specialization and localised competition, European Economic Review, 43, 409-429

and the networks and connections location implies, in fostering innovation and how this can be analysed.

3.   Cities and growth

Cities are agglomerations. They are places in which densities of occupation, of residence, and of employment, are higher than they are elsewhere. It is the process of agglomeration or other kinds of stickiness which is essential to the formation of a city. Simply wishing to be like your neighbours will not create long lived clusters of activity except in a static world[8].

Cities facilitate specialisation and division of labour in a flexible and malleable way. The division of labour in Adam Smith's pin factory was static. The processes involved in making the pin were given and division of labour cut the activity up into its constituent stages, much as Taylorism did for factories in the post war period in America. Almost all those factories are now closed and the firms that owned them have gone out of business or struggle to reinvent themselves. But the cities mostly still exist.

Cities which were based on the factory division of labour, from pin factories to cotton mills to car plants, from Manchester to Detroit, have struggled to escape the static forms innovation. Cities which are based on different forms of specialisation have found it easier to reinvent themselves. London lost 750,000 jobs in manufacturing between 1971 and 2001. But it more than replaced that employment with services based activities. Innovation in communications, in computing, in media, made it possible to create new industries and new specialisations, mixing and matching the skills of the old.

The term agglomeration refers to the way in which activities are stuck together. It gives a good indication as to what goes on in the situations where businesses, people and institutions come together. Even the word itself sounds gluey - and no parts of the economy act in isolation. Isolation is subsistence, where no trade takes place and everything has to be produced by an individual unit. It is the recipe for poverty at best

---

[8] B Rosewell and A Horton, 'Neighbours are not enough: what are the minimum conditions for urban existence' paper presented to Oxford conference on complex systems, date,

and famine at worst. In the UK, the highest output per head is found in the most densely occupied parts of the country. London has the highest productivity at £xx per head employed, and inner London is higher still at £yy per head[9].

Another way to illustrate the importance of density is to compare employment density with earnings. Earnings data is more reliable than output estimates at a local level and provide a good proxy. The chart below shows earnings differentials controlled for sectoral mix and illustrates across Great Britain how this is correlated with employment density.

Figure 3.1: **Earnings differential and log of employment density, 87 GB areas, 2001.**
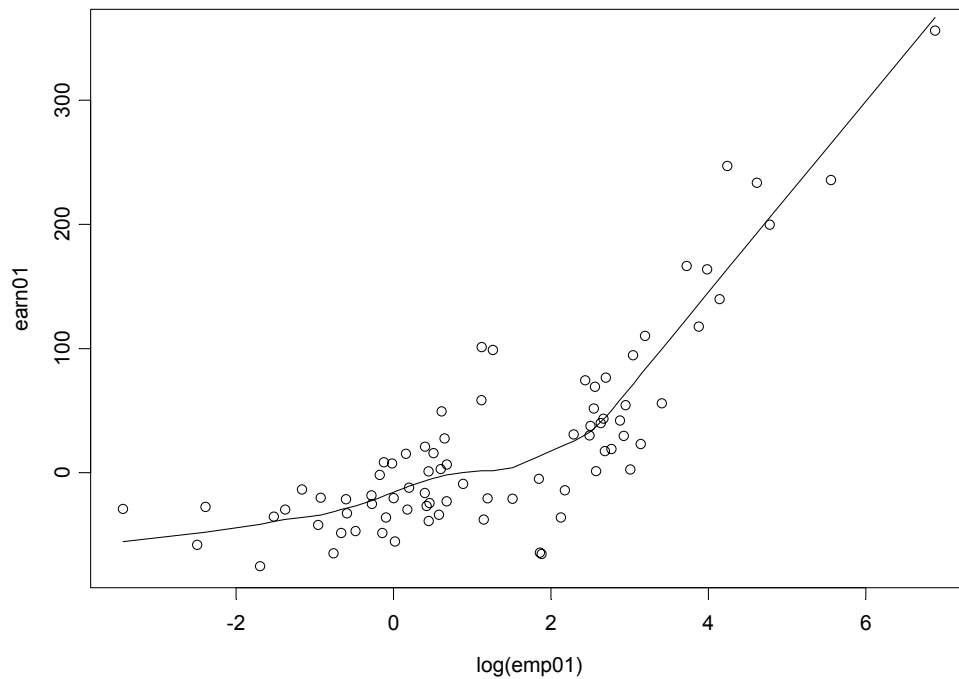
Note: The solid line is the fitted values from the general non-linear local regression technique, with span = 0.6

---

[9] These are the official estimates. They are certainly too low as they assume that headquarters produce no output and that much of financial services is transferred out of London.

Agglomeration gives the opportunity to trade and exchange as well as to build the institutions that regulate and empower such exchanges. It is important to realise that successive developments in technology, which have improved our ability to communicate, have not made cities redundant. Far from spreading out more and more over the available geography, we concentrate in particular areas. The telegraph, trains, telephones, aeroplanes, the Internet and email have all succeeded in creating larger concentrations of activity rather than dispersing us to atomistic locations. We like to travel more, and may engage in video conferencing, but this seems to have intensified our need for face-to-face interactions.

Innovation tends to start at the centre and spread out. The earliest applications of computer technology in business were in banking and accounting records. They required huge installations with major air conditioning. They were tended by men in white coats, providing system capacity which would fit on a PC in today's world. These expensive installations started out in city centres where the systems analysts could look after them and nascent programmers get together after work and swap horror stories. Office blocks in London were filled with highly qualified people trying to work out how to produce accurate bank statements on a monthly basis for the holders of current accounts. All of those jobs and all of those computers have gone again. The computers themselves are nowhere near London and the programmers have moved on to the games industry. This is no longer innovation, it is business as usual. It can happen anywhere.

4. Cities and Networks

We have shown that cities are associated with higher rates of growth, of innovation and with higher levels of productivity. It seems reasonable to hypothesise that this is because of the deeper and denser networks which can be facilitated in cities. This section examines models of the potential development of such networks.

It can be shown that myopic non-rational agents acting in a self-interested way to increase their fitness can build up quite rapidly a densely connected network, even when

they are subjected to regular adverse shocks. Ormerod and Colbaugh[10] consider complex systems whose properties are not static but which evolve dynamically over time.

They consider systems which are populated by agents which are heterogeneous in terms of their fitness for survival. The agents are connected on a network, which evolves over time. In each period, agents take self-interested decisions to increase their fitness for survival to form alliances which increase the connectivity of the network.

The network is subjected to external negative shocks both with respect to the size of the shock and the spatial impact of the shock. Despite this, a densely connected network can evolve rapidly. Further, the overall fitness of the system also rises, despite the fact that the model is deliberately kept as parsimonious and simple as possible, and refrains from incorporating features such as increasing returns and externalities arising from preferential attachment which are undoubtedly a feature of the real world.

The existence of negative shocks, often on a substantial scale, is an important feature of the evolution of cities. Batty[11] shows that although at a given point in time cities scale with size in the upper tails of their distribution, the evident macro-stability in such distributions at different times masks a volatile and often turbulent micro-dynamics, in which objects can change their position or rank-order rapidly.

Initially, we have a model populated by N autonomous agents. These are placed on a circle, with the location of each agent drawn from a uniform distribution. The k nearest neighbours of each agent are therefore defined unequivocally. In this context, it is important to note that the phrase 'nearest neighbours' means nearest in terms of a particular attribute, parameterised by their location on a circle, so the attribute could be 'geography', 'industrial complementarity', or whatever.

Each agent is assigned a fitness level, f_i, chosen at random from a uniform distribution on [0,1]. The model evolves in a series of steps over time. In each step, the model is

---

[10] P Ormerod and R Colbaugh, 'Cascades of Failure and Extinction in Complex Evolving Networks' , *Journal of Artificial Societies and Social Simulation*, 4, 2006
[11] M Batty, 'Rank clocks', *Nature,* Vol 444|30 November 2006| doi:10.1038/nature05302

subjected to an external shock. The size of the shock, q_j, is drawn in each period from a random distribution bounded in [0,1]. An agent is selected at random to be the location where the shock hits the network. The spatial impact of the shock, s_j, is drawn from a random distribution also on [0,1]. All agents which are within the distance s_j of the agent where the shock hits also receive a shock of the same size.

The fitness of shocked agents is decreased by the size of the shock. Agents whose fitness level $f_i < q_j$ are deemed to become extinct. So extinctions apply directly to nodes: the node (agent) receiving the shock and all nodes (agents) within the distance of the shock receive it.

An extinction event of size m is defined as one in which the proportion m of all agents becomes extinct. In the next period, an extinct agent is replaced immediately by an agent with the identical fitness f_i. Note that a replacement rule in which the new agent has a fitness chosen at random from, say, a uniform distribution is more likely to increase the overall fitness of the system. On average, agents which become extinct will have low fitness levels, and so if they are replaced by agents with random fitness, the overall fitness of the system is likely to rise. This rule therefore eliminates this bias towards increasing fitness.

The model proceeds on a step by step basis, and in each step each pair of agents can form an alliance with fixed probability p. If the alliance goes ahead, the fitness of the each agent is increased. The new fitness is given by

$$f_{i,j} \; = \; f_{i,j-1} + v_{ij} \; - \; f_{i,j-1} * v_{ij} \qquad\qquad (1)$$

where $f_{i,j}$ is the fitness of the i th agent in period j and $f_{i,j-1}$ its fitness in the previous period, and the v_{ij} are drawn from a uniform distribution on [0,1]. The expression (1) ensures that the fitness of each agent is bounded in [0,1], and implies that the value of alliances is subject to diminishing returns. The value of each alliance is defined as ($f_{i,j}$ - $f_{i,j-1}$).

The fitness of agents therefore rises as the connectivity of the graph increases. However, an agent with an alliance to another agent will also receive any shock received by the latter, even if the agent is beyond the physical distance s_j of the shock. So the capacity of shocks to spread spatially is increased. The reduction in fitness transmitted to an agent is $q_j*(f_{i,j} - f_{i,j-1})$,, in other words the size of the external shock received by an agent multiplied by the value of the alliance. The shock is transmitted across any sequence of alliances in the network, until the reduction in fitness at the relevant step in the sequence falls below b (where b is small).

Three statistical distributions are used for both the size and the spatial impact of the shocks, to check the robustness of the results: a uniform on [0,1], a normal with mean = 0.5 and standard deviation = 0.1, and a beta with parameters $z_1 = 1$ and $z_2 = 5$.

The connectivity of the graph which evolves as alliances are permitted obviously depends upon p, the probability of agents forming alliances in any step of the solution. However, in any individual solution, the evolution of connectivity over time is by no means smooth because the alliances of agents which become extinct disappear.

It is important to recall that this model differs considerably from almost all of the literature which considers cascades of information/failure/extinction across networks of agents. These in general consider networks of a given type whose structure has already evolved. So it is useful to consider the type of graph which emerges as the networks evolve due to decisions of self-interested agents and external negative shocks.

In step 1 of each solution, each agent pair (1 and 2, 1 and 3…1 and N, 2 and 3,…..) has a connection formed with probability p. At this point, by definition the degree distribution is that of a standard random graph. A shock is then applied, which might result in some agents becoming extinct. They are replaced by other agents. In step 2, each agent pair which is not already connected has a connection formed, again with probability p. So in the early steps of any particular solution of the model, for low values of p and for low values of shocks both in size and range, the degree distribution continues to approximate a random graph. Obviously, in the limit if there were no shocks, the network would become completely connected.

Beyond the initial steps of a solution, it is not really possible to characterize precisely the nature of the graph because it is always evolving. New connections are being formed in any given period, and existing ones are being removed as agents become extinct. The degree distribution of the graph never settles to an equilibrium.

However, the network appears to have a tendency to fluctuate between periods when the degree distribution is right-skewed and periods when it is more similar to a random graph.

An illustration of this is given in Figure 4, which shows the evolution of the size of the principal component of the graph in the initial 40 steps of an individual solution of the model. This illustrates the degree of connectedness shown at each step of a typical solution for different ways of drawing the shocks and different probabilities of forming an alliance. Thus, for example, in the case of shocks drawn from a uniform distribution and a probability of forming alliances of 1 per cent for each pair, the degree of connectedness varies quite substantially, while this is less the case if the shocks are drawn from a beta distribution.

Figure 4.1   *Size of principal component of graph in first 40 steps of a typical solution of the model. Beta refers to the size of the shock drawn from a beta distribution and its range from a normal; normal and uniform refer to both size and range drawn from a normal and a uniform respectively; pAlly is the probability of forming an alliance (p)*
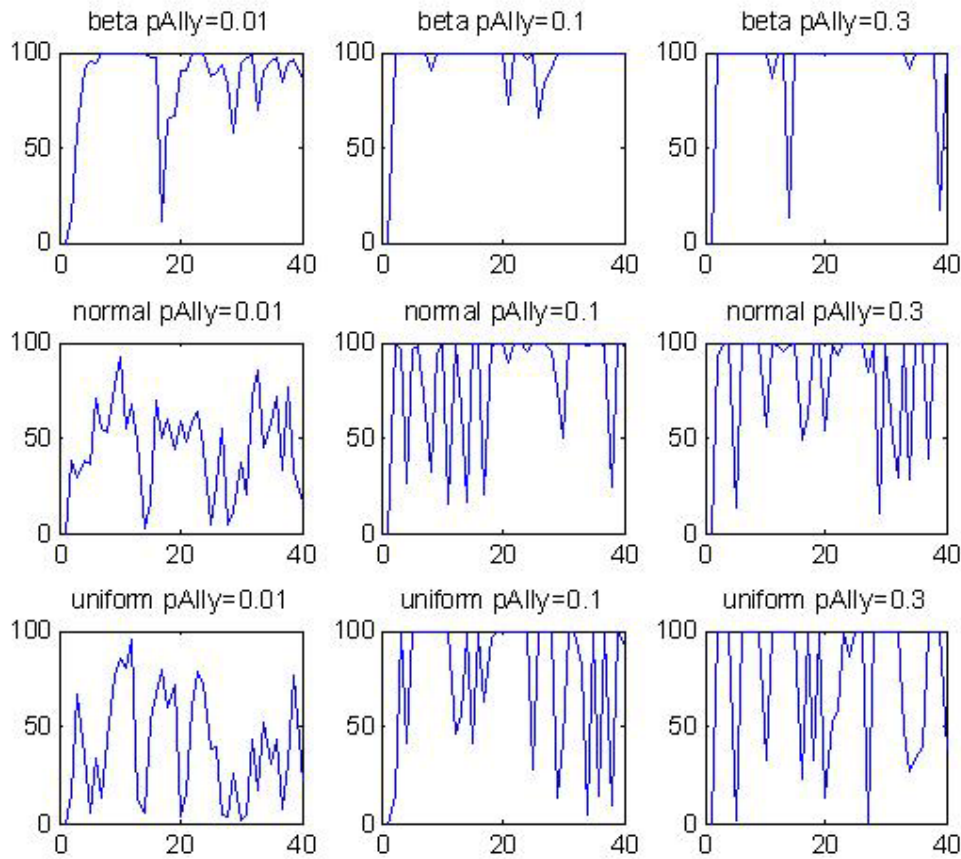
It is clear that even in the presence of persistent and non-trivial shocks, a densely connected network of agents can be built up quite rapidly and also varies substantially. In the context of cities, it is therefore reasonable to think of each time step in the model as corresponding to a substantial period in real time. We have not attempted a precise calibration, but we might usefully think of each period in the model being of the order of a decade in real time.

Cities are locations which facilitate the ability to make connections and form alliances. The network formation described in this section and its ability to evolve is

a good potential description for city networks. The ability to survive shocks and to remain connected is also relevant, since cities have traditionally had a defensive and protective role as well as a role in fostering performance

The next step is therefore to consider diffusion across the topology of the system at any point in time. In other words, we hold the evolving topology at a point in time fixed, and consider how innovation might spread.

5.   Networks and Innovation

The role of networks in enabling the diffusion of innovation has received enormous attention. From the identification of S-curves and the move from early adopters to majority take-up, to the application of small world scale free networks the issue of diffusion has attracted much research. Duncan Watts (2003)[12] has started to apply his analysis of networks to business issues, largely however focusing on problems of negative shocks. For example, he looks at how the Toyota network of firms were able to respond in a creative way to an apparently disastrous fire in one supplier's factory and thereby restart production by ad hoc adjustments.

Here, we take a particular topology from the evolving network. Agents can be in one of two states of the world, 0 or 1, corresponding to whether or not they have adopted a new innovation. Each agent is assigned a threshold value, τ_i, drawn at random from a uniform distribution on [0,1]. Agents will switch from state 0 to state 1 if the proportion of all agents to which they are connected is $> \tau_i$ otherwise they remain in state 0.

In other words, the model is an example of the class of models described by Schelling as 'binary decisions with externalities'. An agent chooses between two alternatives, whether or not to adopt the innovation, and this decision has consequences for other agents. The greater the proportion of agents which has already adopted the innovation, the more likely it is that the next agent who considers the issue will also decide to adopt.

---

[12] D J Watts, Six Degrees, Science for a Connected Age, Heinemann, 2003

Watts[13] has considered a model very similar to this using random networks with different degreees of connectivity. Initially, all agents are in state 0. A small number are selected at random to move to state 1. The model then proceeds on a step-by-step basis, and in each step, each agent who is not already in state1 is called to decide whether or not to move to state 1 or to stay in state 0. The process continues until an equilibrium distribution of agents in states 0 and 1 obtains, defined as being either when all agents reach state 1, or when no agents switch during a particular period.

The model is solved many times, using the same network, but with different agents[14] called at random to move into state 1. We observe the extent of the cascade across the network, in other words the proportion of agents who adopt the innovation.

In this model, only agents who have alliances in the evolving network model described above are put into the cascade process, so in the more sparsely connected networks which are generated, there will be a limit to the process of diffusion. Not all N agents can adopt the technology, because some will have no connections ('alliances').

We choose examples from solutions of the evolving network model in which the size and range of the external shocks is drawn from a XX distribution and the probability of forming an alliance is 0.01.

Initially, 3 out of the 100 agents are selected to move into state 1. In the most connected of the networks we considered, 97 out of the 100 agents have alliances, and the average number is 9.8.

Figure 5.1 shows the results.

---

[13] DJ Watts, 'A simple model of global cascades on random networks', *Proceedings of the national Academy of Science,* 5776-5771, 2002

[14] Of course, there is a very small probability that in two different solutions of the model, these agents chosen at random are identical in the two solutions
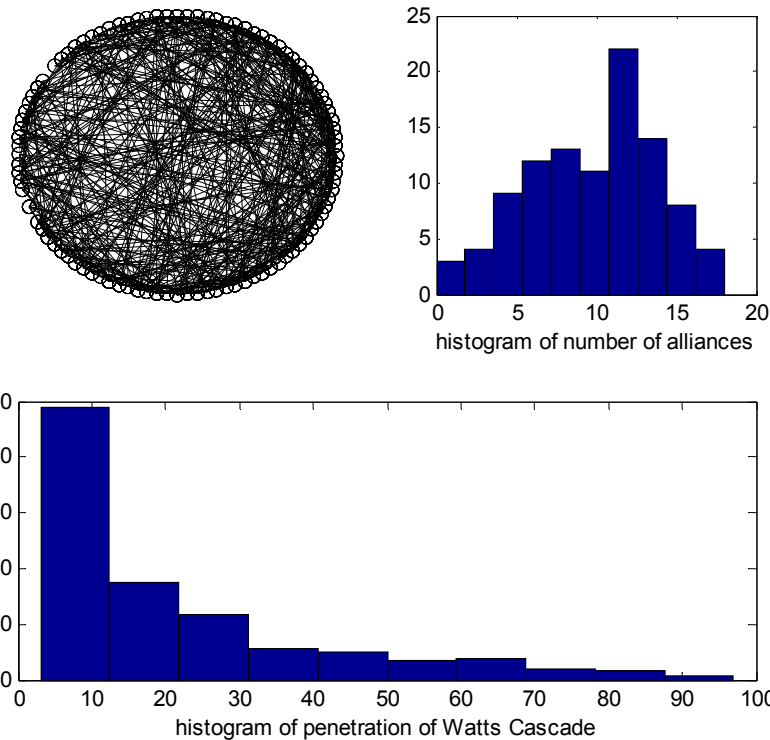
Figure 5.1 *Network in which 97 per cent of the agents are connected, average number of connections across all agents is 9.8. Top left chart plots the network, top right the distribution of the number of connections per agent. The bottom chart shows the outcome of the cascade process, the number of agents who switch to state 1 using a binary decision rule with heterogeneous thresholds, initially 3 agents are selected to move to state 1; 1000 separate solutions*

As is clear, the network here is densely connected. Even so, it is only in a distinct minority of cases that the innovation diffuses across a majority of the agents. This seems realistic, since we know that most innovations do not succeed in becoming adopted generally (although this model does not take into account any potential competitive benefits which adoption of the innovation might bring, which would alter endogenously the incentives to adopt).

The propagation of the innovation in this densely connected network is limited by the relatively large number of connections which each agent has. The larger the number, the less likely it is that the threshold proportion for any given agent will be exceeded.

Figure 5.2 shows the results for the period immediately after the network used in Figure 3, when the evolving model had received a large external shock. 64 of the agents now have alliances, with an average of 5.4 per agent.
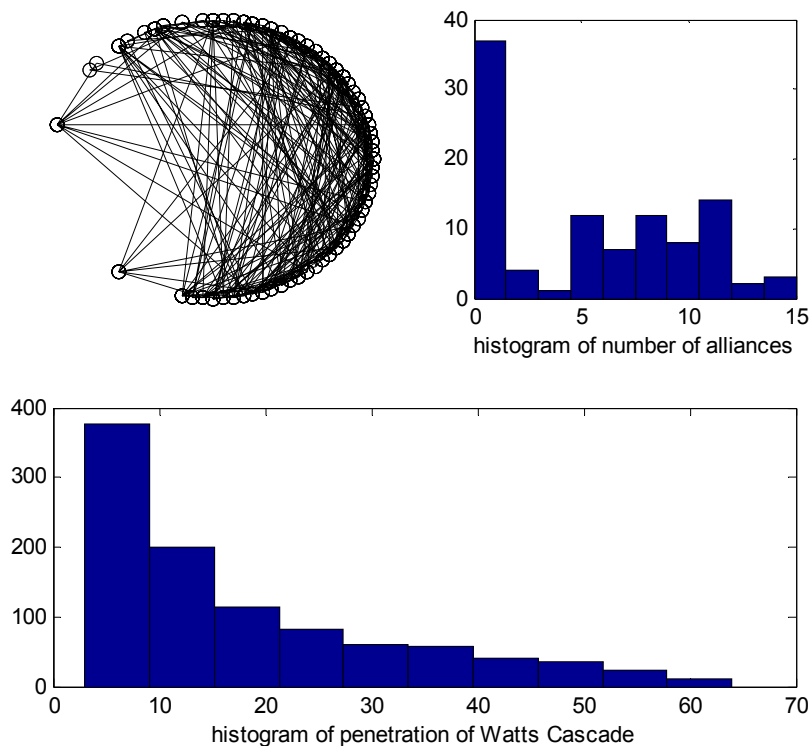

Figure 5.2 *Network in which 64 per cent of the agents are connected, average number of connections across all agents is 5.4. Top left chart plots the network, top right the distribution of the number of connections per agent. The bottom chart shows the outcome of the cascade process, the number of agents who switch to state 1 using a binary decision rule with heterogeneous thresholds, initially 3 agents are selected to move to state 1; 1000 separate solutions*

Here, the diffusion process is more limited because by definition 36 per cent rather than just 3 per cent of agents are excluded from it. Even so, the general shape of the cascade distribution is qualitatively similar to that of Figure 3, with diffusion occasionally reaching across all eligible agents (64 per cent).

Finally, in Figure 5.3 we show diffusion in a weakly connected network, in which only 40 per cent of agents are connected, with an average number of connections (across all N agents) of 0.9.
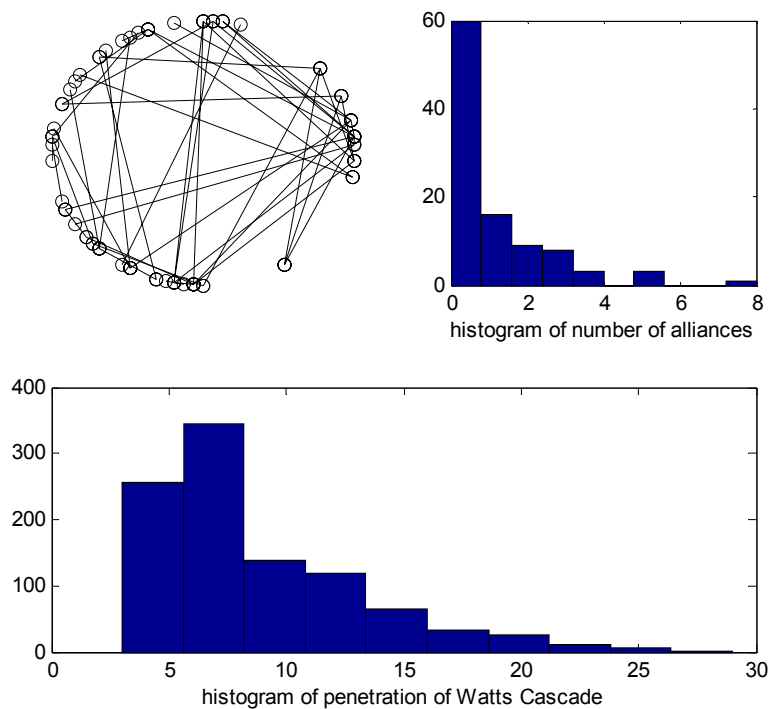

Figure 5.3 *Network in which 40 per cent of the agents are connected, average number of connections across all agents is 0.9. Top left chart plots the network, top right the distribution of the number of connections per agent. The bottom chart shows the outcome of the cascade process, the number of agents who switch to state 1 using a binary decision rule with heterogeneous thresholds, initially 3 agents are selected to move to state 1; 1000 separate solutions*

In this example from the evolving network, diffusion is limited both by the proportion of agents with no connections, and by the weak connectivity of those which are connected. Even amongst those agents with connections, the percentage who switch to state 1 is never more than 70 per cent.

So we do not need a completely connected network to observe widespread adoption of an innovation. Figure 5.3 shows that even with average connections per agent being less than 10 per cent of the total possible, cascades of adoption can arise on a near global scale. However, in networks which are considerably less well connected, diffusion is much more sharply limited, both by the proportion of agents operating in isolation, and by the difficulty of diffusion across the weakly connected network of the agents who do actually have links with others.

6.   Conclusion

Cities are productive and growing places. Understanding how networks form in cities and hence the role of such networks in ensuring that innovations can be produced and diffused is clearly important and under-researched.

The ability of dense locations to produce more growth appears to be well established and in the UK, the available data suggests it may even be strengthening. Moreover, there is so far no evidence that allows an estimate of the density at which agglomeration effects would tail off and crowding costs would outweigh such benefits.

Cities are agglomerations and such agglomerations tend to exhibit higher productivity than more dispersed locations. Productivity and economic growth are themselves driven in the long term by innovation and its diffusion. Cities make innovation easier to diffuse through network effects, competition and straightforward example. The fact that cities grow fastest when economies are also growing is prima facie evidence for this.

This paper has explored some of the network characterisations that might be most relevant to cities. In particular we give an example of an evolving network in which alliances are made and maintained in the face of negative shocks.

We then examine the ability of individual states of these networks to support effective diffusion of innovations. The results suggest that widespread adoption of innovation is supported even where the average number of connections that each agent has is quire small.