Network models of innovation process and policy implications[1]

Paul Ormerod, Bridget Rosewell,
Greg Wiltshire

Volterra Consulting

October 2009

Corresponding author brosewell@volterra.co.uk

[1] We are grateful for the support of the Manchester Independent Economic Review in undertaking much of this research

Abstract

Innovation is clearly a disequilibrium phenomenon and therefore a consideration of the path that innovation takes through the economic system is crucial to understanding the adjustment process for the economy. We have developed models of how innovation can pass through the economy calibrated to the actual relationships of businesses in different industries. These models consider not only the network characteristics but also the knowledge and competence of firms and their willingness to pass on new ideas. All of the results are calibrated to specially commissioned survey results. We consider the steps required to generate an innovation cascade and the time likely to be required to achieve this. Further, we review whether current policy programmes have the potential to be effective in fostering the effective introduction of innovation in the context of these data and models.

1 Introduction

Understanding the innovation process is key to understanding how capitalism has created levels of productivity which were previously undreamed of. However, the standard models of the firm or of macroeconomic growth do not address this very effectively. This is largely because innovation is fundamentally a disequilibrium phenomenon – it is about the constant adjustment towards a new equilibrium that is never reached because further innovation is always occurring. Innovation is a disturber and a re-adjuster, to both the firm and the economy. Since innovation is a dynamic process analysis using comparative statics fails to capture the nature of the process. Schumpeter pointed out long ago that *"innovation…does not lend itself to description in terms of a theory of equilibrium"* (Schumpeter, 1928)[2], and more recently Antonelli notes that *"innovation is the distinctive element of a dynamic process which cannot be analysed with the equilibrium approach"* (Antonelli, 2008)[3].

This paper addresses the dynamics of innovation by considering the process by which a particular innovation might be made effective. To do this, we look at the networks by which organisations might discover and adopt innovations. Glückler notes that growth

[2] Schumpeter, J. A. (1928) The instability of capitalism, *Economic Journal*

[3] Antonelli (2008) *Localised Technological Change: Towards the economics of complexity*, Routledge

and innovation largely result from network dynamics, and this is the starting point this study takes (Glückler, 2007)[4].

Innovation is the deliverable realisation of an invention. It is important not to treat technological improvement or invention as identical to innovation (Metcalfe 2007)[5]. Metcalfe notes that even were technological advancement to end today, there would still remain a considerable amount of potential for innovation within an economy. Innovation is more than just invention; it can also involve the use of a better production process, a new service, or the changing of the network structure itself (Antonelli, 2008). Metcalfe tells us "*innovation requires access to and command of many more kinds of knowledge and capability that are summed up by the phrase 'science and technology'*". Innovation is the economic realisation of an invention and requires an understanding of consumer and user needs as well as knowledge of the market and organisation if it is to occur. Primarily it is the application of new combinations of resources to the economic process.

The most relevant indicators of innovative outcomes should reflect the efficiency and effectiveness in producing, diffusing and exploiting economically useful knowledge (Lundvall, 1992)[6]. One measure of the success of an innovation is the proportion of the total possible market the innovation has penetrated (Mansfield, 1961)[7]. We model the spread of an innovation through the market and, through seeing how far the innovation has penetrated, gain a metric that measures innovation.

We go on to look at the impact that different structures have on the speed and effectiveness of the spread of innovation and how public policy might affect this.

A network model of the innovation process needs to capture those features that are relevant to the spread of innovation between firms. We estimate the number of links

---

[4] Glückler, J. (2007) Economic geography and the evolution of networks, *Journal of Economic Geography*

[5] Metcalfe, J. S. 2007 lecture delivered at University of Jena

[6] Lundvall, B. A. (1992) Explaining Inter-firm Cooperation and Innovation--Limits of the Transaction Cost Approach, in: Grabher G (ed.) *The Embedded Firm: On the Socioeconomics of Industrial Networks,* Routledge, London.

[7] Mansfield, E. (1961) Technical Change and the Rate of Imitation, *Econometrica*

firms have with others, their own ability to generate innovations, and their abilities to both disseminate and absorb knowledge and innovation from the wider network of firms in which they are embedded.

The ability of an industry to have a high degree of innovation will be a function of various features of the network and the firms within it. The number of links within a network may have a positive or negative bearing on its amount of innovation (Ormerod and Rosewell, 2008)[8] depending on the decision rule used by agents in deciding whether or not to adopt the innovation. The propensity individual firms have to innovate, disseminate innovation and absorb innovation will all be important determinants of an industry's innovation.

We use data from a study of Manchester in the UK to calibrate a model of this innovation process.

2 The Network Model

2.1 Methodology

Figure 2.1 outlines the approach used to model the adoption of a single innovation across a network while Table 2.1 describes the parameter values and ranges used in this study.

Figure 2.1: Model overview

[8] Ormerod, P. and Rosewell, B. (2008) Innovation, Diffusion and Agglomeration, forthcoming in *Economics of Innovation and New Technology*

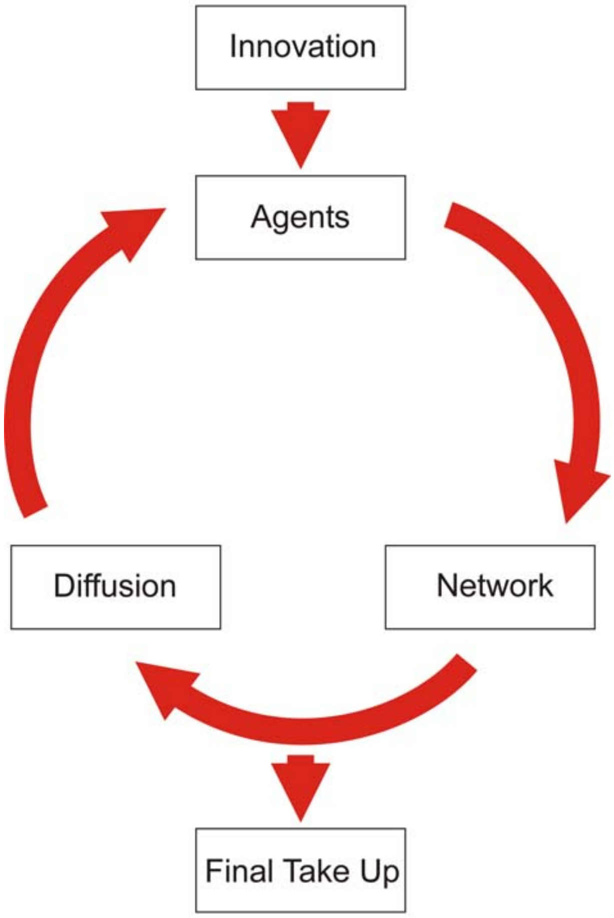

The model takes an initial innovation to be exogenous and it is taken up by one agent/organisation at the outset. The characteristics of the agents are governed by their willingness to innovate, their desire to keep innovation to themselves, and their willingness to communicate with others. The innovating agent will be connected to other agents via the network structure, and at the next step of the model the innovation will be passed on according to the extent agents discover the innovation and their own willingness to take it up. At further steps of the model further agents may be able to discover and take up the innovation, until eventually no further take up occurs.

2.2 Innovation Behaviour of Organisations

We define two different methods by which innovation may be passed on via the network linkages. The first is a direct relationship between two partners, while the second is a group relationship.

First, an organisation with an innovation will provide it to another firm only if its level of **secrecy**, or the propensity of a firm to try to retain the benefits of its innovations, is less than the **absorptive capacity,** or the degree to which a firm actively engages in activities which enable it to identify and adopt new innovations, of the firms it is linked with. This method of adopting an innovation represents a mutual relationship or ***exchange*** between firms and implies a degree of trust or collaboration. It is probable that this relationship is more likely to exist with customers, suppliers or third parties than with competitors.

The second method for spreading an innovation we describe as a ***copying*** behaviour. Here if a firm looks at the spectrum of organisations to which it is linked and finds that the proportion that have adopted an innovation is higher than their own personal threshold, they will mimic their behaviour and adopt the innovation. In some circumstances this threshold may be very high and only when all or nearly all of the firms an organisation has relationships with have taken up an innovation will they be persuaded to do the same. For other organisations relatively few businesses may have to have the same innovation before they adopt it. This mechanism represents a copying behaviour. This may occur even when a firm may not fully understand the reasons and benefits of an innovation but relies on observing that other businesses have adopted it. This behaviour is more likely to be a response to competitor behaviour.

There is a growing body of literature that indicates that a range of macro consumer behaviour, such as the take up of bank accounts by unemployed people[9] and the binge drinking behaviour of young adults in the UK[10], can be explained by such decision making mechanisms.

2.3 Network Behaviour

---

[9] Volterra Consulting Ltd (2006) The Spread and Containment of Behaviour Across Social Networks, prepared for the Financial Services Authority, UK

[10] Ormerod, P. & Wiltshire, G. (2008) 'Binge' drinking in the UK: a social network phenomenon; and Nature News (23 Jun 2008), doi: 10.1038/news.2008.910, News

The network behaviour is not set *a priori*. One of the important aspects of this paper is that we are able to calibrate the network across which any innovation percolates to the actual networks that we observe in our survey. To enable us to use these distributions in the model a variable is introduced called the network scaling factor (S_0). This is essentially the maximum number of links, or degree, a firm can have in the model (while the minimum is ⅛S_0). Once this is set the number of links of all firms can be determined. As S_0 is increased the number and therefore the density of links is increased.

The *k*-clique Algorithm

As discussed the model starts by taking the value for S_0 and calculating the implied degree distribution of links. It then builds the network to match this distribution as closely as possible by generating a series of what are called k-cliques. For example, there may be three firms who should have two links, four firms that should have three and five firms that should have four links. In this case the firms that require the same number of links are connected into a fully connected sub-graph or k-clique as shown in Figure 21. Links formed using this method are reciprocated, which means that an innovation can be spread in both directions.

Figure 2.2: The initial stage of generating a network structure for a given industry using k-cliques.

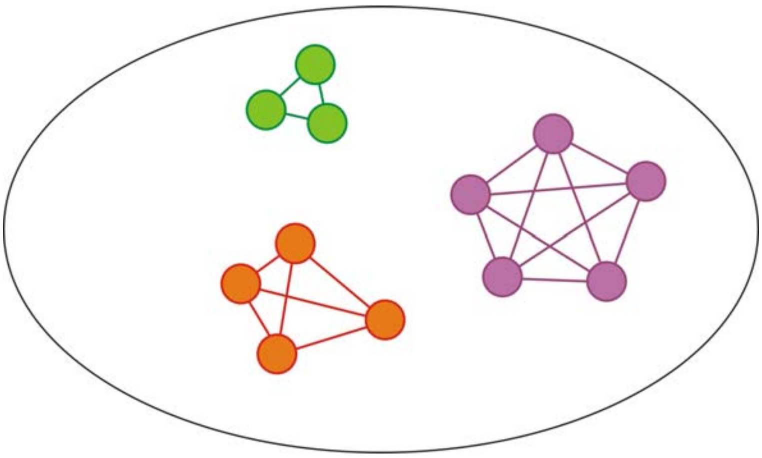

At this stage the network is therefore highly clustered and each cluster is isolated. In order to move towards a more realistic network structure half of the links in the model are re-wired to randomly selected firms while the degree of each firm is retained, as shown in Figure 2.2. What is achieved is a hybrid network structure where quasi-clusters of firms exist as part of a larger region wide network. However, it is not always possible to guarantee that the program will find a way to rewire 50 per cent of the links within the computational constraints, Table 3.2 presents the average percentage that the algorithm achieved across the 1000 networks generated for each industry parameterisation.

Figure 2.3: The second stage of generating a network structure for a given industry where the k-cliques are rewired but the degree distribution is retained.

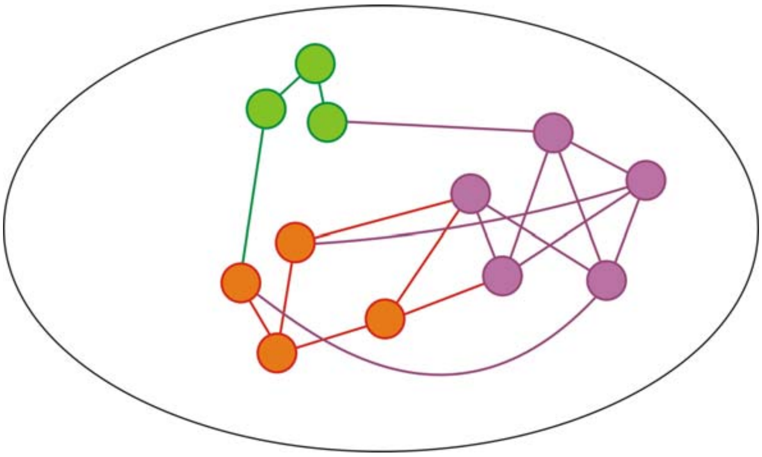

2.4 The Parameters

The model parameters are set out formally below, alongside the ranges used in the model, which reflect the data used and are more fully described in Section 3.

Table 2.1 Parameters used in the generation of the network model for each industry

Parameter	Description	Value/ Range
n	Number of agents or firms in network.	1000
α	The lower limit to the distribution for the threshold of agents to adopt an innovation based on an evaluation of agents connected to them by their business network. – copying	0-1.1
ε	The upper limit to the distribution for the threshold of agents to adopt an innovation based on an evaluation of agents connected to them by their business network. – copying	1-1.1
S_0	Network scaling factor, equals the highest number of business links a firm can have in the model. – size of network	8-64
β_a	Absorptive Capacity Scalar, this is the scaling factor applied to each firm's absorptive index when drawn from the industry distribution	0-10.05

β_s	Secrecy Scalar, this is the scaling factor applied to each firm's secrecy index when drawn from the industry distribution	1
periods	The number of time steps from the point that the first innovation is generated, that the model runs for.	50
Repeats	The number of iterations that each parameterisation is repeated in order to find by the average and range of possible outcomes.	1000

The steps that the model takes in disseminating an individual innovation are described in detail in the flow charts below. The model enables us to determine the final take up of an innovation, given the parameters of the model. It also allows us to consider the effectiveness of different behaviours, agent characteristics and network structures on the likelihood of an innovation being adopted by all firms – a global cascade.

Figure 2.4: The network model

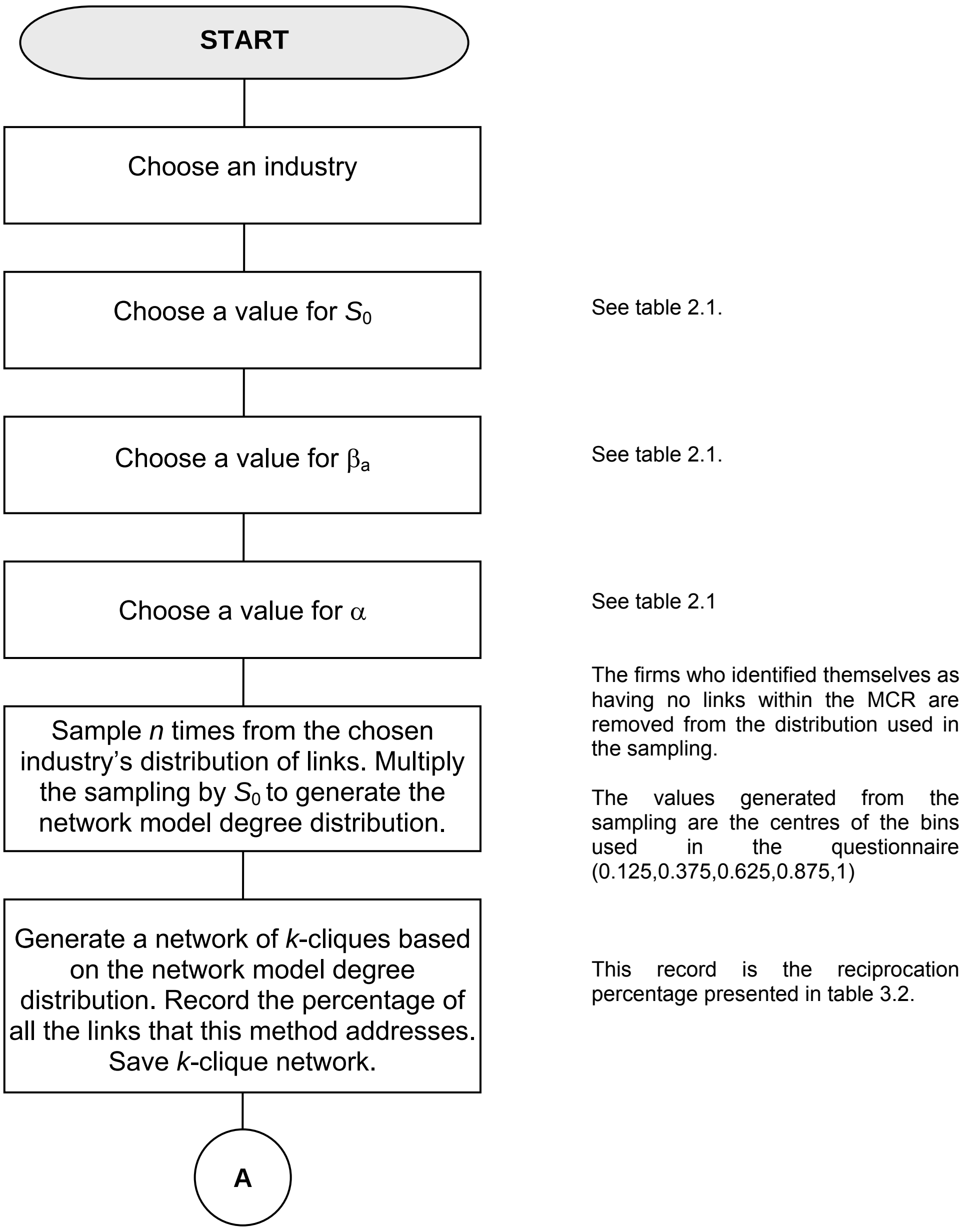

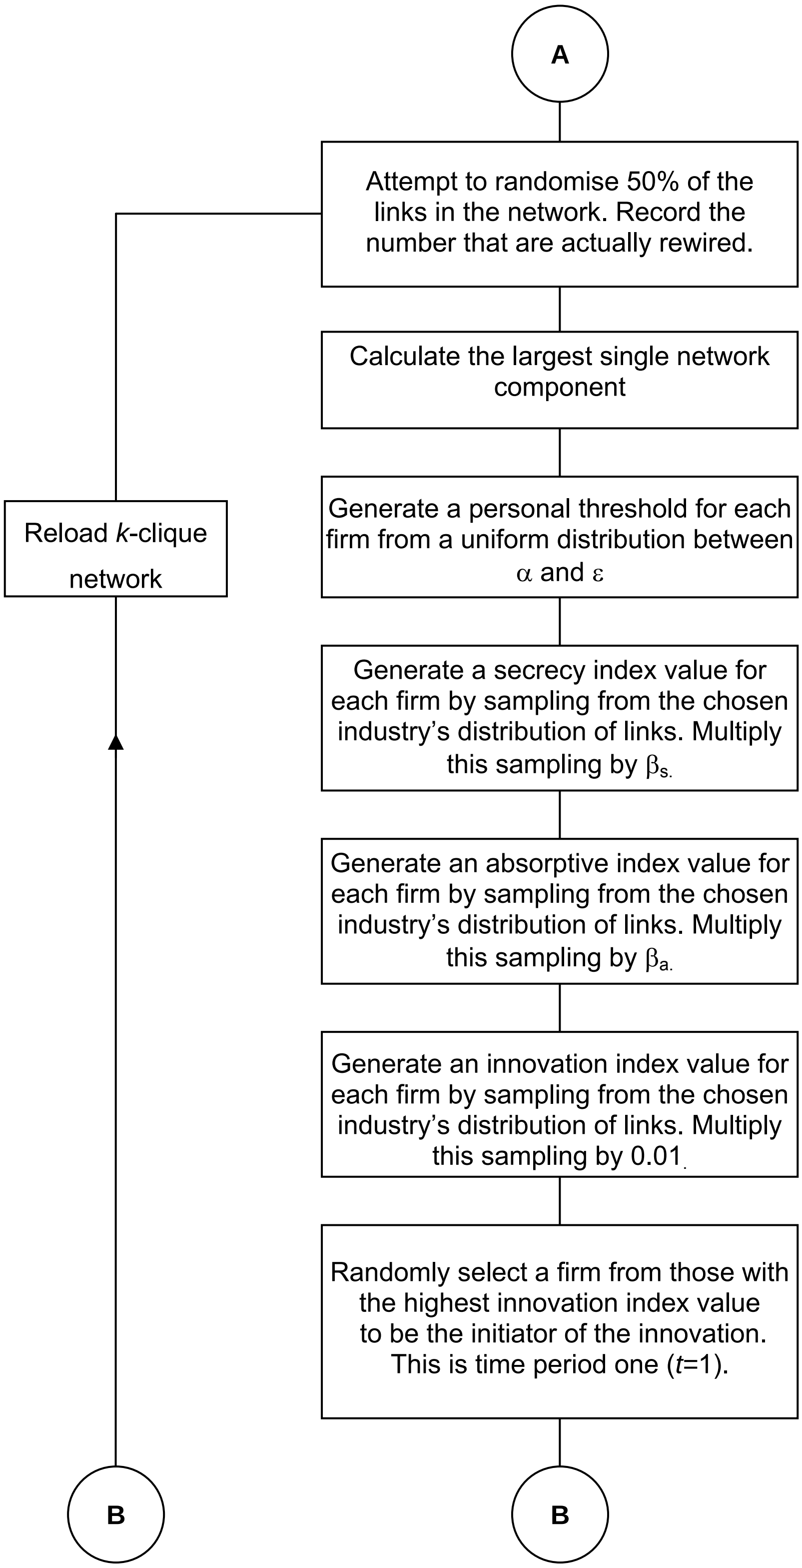

This record is the randomisation percentage presented in table 3.2.

This figure is the super-component percentage presented in table 3.2.

See table 2.1

The values generated from the sampling are the upper limits of the bins (0, 0.1, 0.2 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0).

The values generated from the sampling are the upper limits of the bins (0, 0.1, 0.2 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0).

The values generated from the sampling are the upper limits of the bins (0, 0.1, 0.2 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0).

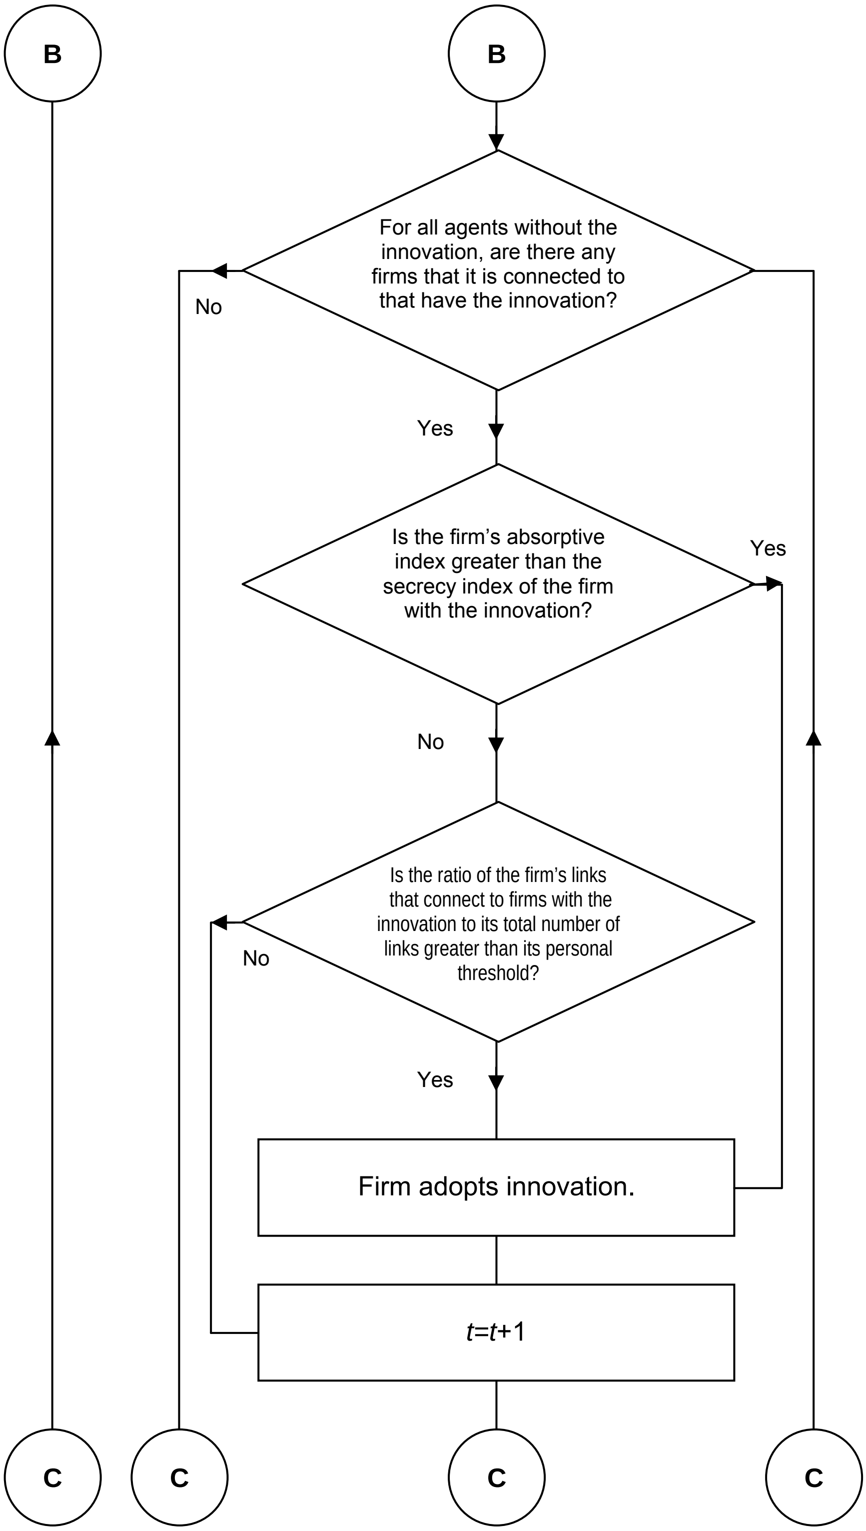
B
B
For all agents without the innovation, are there any firms that it is connected to that have the innovation?
No
Yes
Is the firm's absorptive index greater than the secrecy index of the firm with the innovation?
Yes
No
Is the ratio of the firm's links that connect to firms with the innovation to its total number of links greater than its personal threshold?
No
Yes
Firm adopts innovation.
t=t+1
C
C
C
C

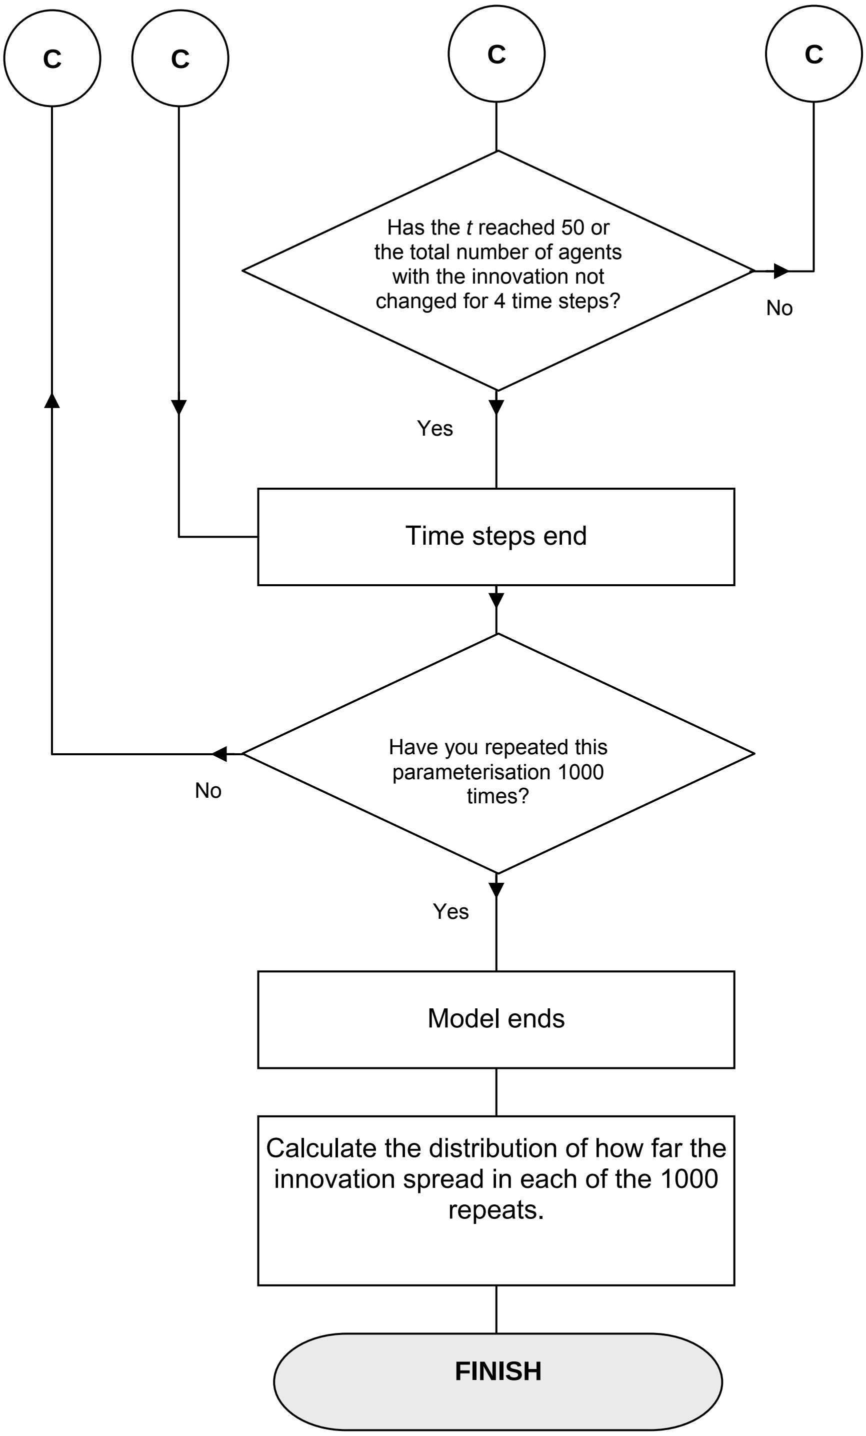
C
C
C
C
Has the t reached 50 or the total number of agents with the innovation not changed for 4 time steps?
No
Yes
Time steps end
Have you repeated this parameterisation 1000 times?
No
Yes
Model ends
Calculate the distribution of how far the innovation spread in each of the 1000 repeats.
FINISH

2.5 Choice of Model Parameter Ranges

Number of agents - n

The number of agents is chosen as a balance between providing a population of firms that enable the various distributions to be well reflected in a single repeat of the model and the computational demands of running 315 model parameterisations, for each of four industries, all of which are averaged 1,000 times. Even at a second per repeat this equates to 15 days of pure computation, a time that does not include the time needed to build the 20,000 network structures.

Boundaries to copying - α and ε

α and ε are used to define the maximum and minimum boundaries of the uniform distribution used to allocate firms their personal threshold used in their copying behaviour. The maximum is held at one throughout this study while the minimum varies. The impact of this is shown in Figure 2.3, as α is increased, on average, firms require a higher proportion of their connections to possess an innovation before they will adopt it. The values of α in Table 8 studied here relate to the entire range of possibilities between 0 and 1. The special case of α and ε equal to 1.1 corresponds to the copying behaviour being turned off.

No data is available from our survey to enable us to calibrate this kind of behaviour.

Figure 2.3: Specification of α *and* ε *for determining the distribution used for the personal threshold of agents used in their mimicking behaviour for (a)* $\alpha=0$ *and* $\varepsilon=1$ *and (b)* $\alpha=0.5$ *and* $\varepsilon=1$.

(a)

(b)

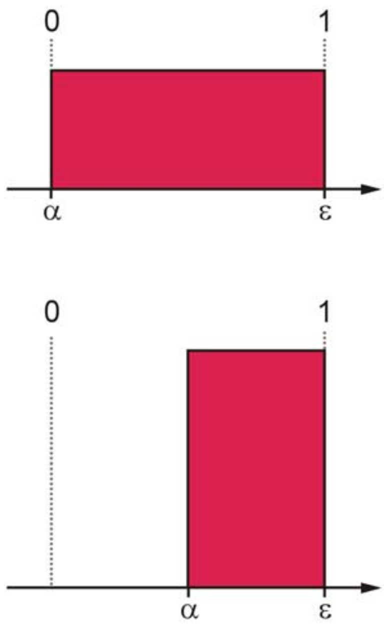

Maximum number of links S_0

S_0 is the maximum number of links of any firm in a given network structure that is then used to define the degree distribution for all other firms. The distribution used is based on the survey data discussed in Section 3.1, where the smallest non zero bin is 0.125. In order to avoid biasing the model by rounding the number of links a firm has it is therefore necessary that S_0 is an integer multiple of eight. It follows that eight itself is the lowest possible value as this will produce firms with only a single link in the network (8x0.125=1). The upper range of S_0 is open ended: for the purposes of this study connectivity an order of magnitude higher than that implied by the minimum was studied.

The survey can only tell us the distribution of links and not the total number, so this is an essential assumption. Results are later shown for different configurations of S_0.

Scaling factors on absorption and secrecy - β_a & β_s

The parameters β_a and β_s are used to vary the values, but not the distribution of, the absorptive and secrecy indices of the firms. When these are initially drawn the values generated are the upper limits of the bins used in the survey (0, 0.1, 0.2 …1.0), these are then multiplied by the values of β_a and β_s. Figure 2.4(a) shows this for the generalized case, for example when β_a equals β_s a firm with an absorption index of $0.5\beta_a$ (the red area) will be able to adopt all of the innovations of firms with an secrecy index of up and including $0.4\beta_s$ (the green area). Figure 2.4(b) and (c) show two examples where the two scaling factors are unequal. When β_a =0.25 and β_s =1 the same firm will only absorb the innovations of firms with secrecy indexes of 0.1 and zero whereas when β_a =2.05 and β_s =1 the same firm will be able to absorb the innovations of all firms.

These examples demonstrate three characteristics of the engineering of this mechanism and the parameter ranges presented in Table 2.1. First, that this research does not have to explore a parameter space for β_a and β_s independently. This is because increasing β_a is equivalent to reducing β_s, therefore we hold β_s=1 throughout this study and only modify the characteristics of the exchange behaviour by changing the value of β_a. Second, it explains why the model requires firms to have a *higher* absorptive index than the secrecy index of a firm that has an innovation and that it is connected to, rather than *higher or equal to*. We do this to avoid the case that a firm with an absorptive index of zero could adopt the innovation of a firm with a secrecy index of zero, which would be counter intuitive, especially when β_a=0. The consequence of this is that when β_s=1 and β_a=0 the exchange method of innovation adoption is turned off and that when β_a>10 all the firms, except those with an absorption index of zero, adopt any innovation they are connected to because any non-zero absorption index will be larger than any secrecy index..

Figure 2.4: Specification of the absorptive capacity scalar (β_a) *and secrecy capacity scalar* (β_s) *to change the amount of adoption of innovation, due to the exchange mechanism, available in a model for (a) the generalised case, (b) the case when* β_a=0.25 *and* $\beta_{s=1}$ *(c) the case when* β_a=2.05 *and* β_s=1.

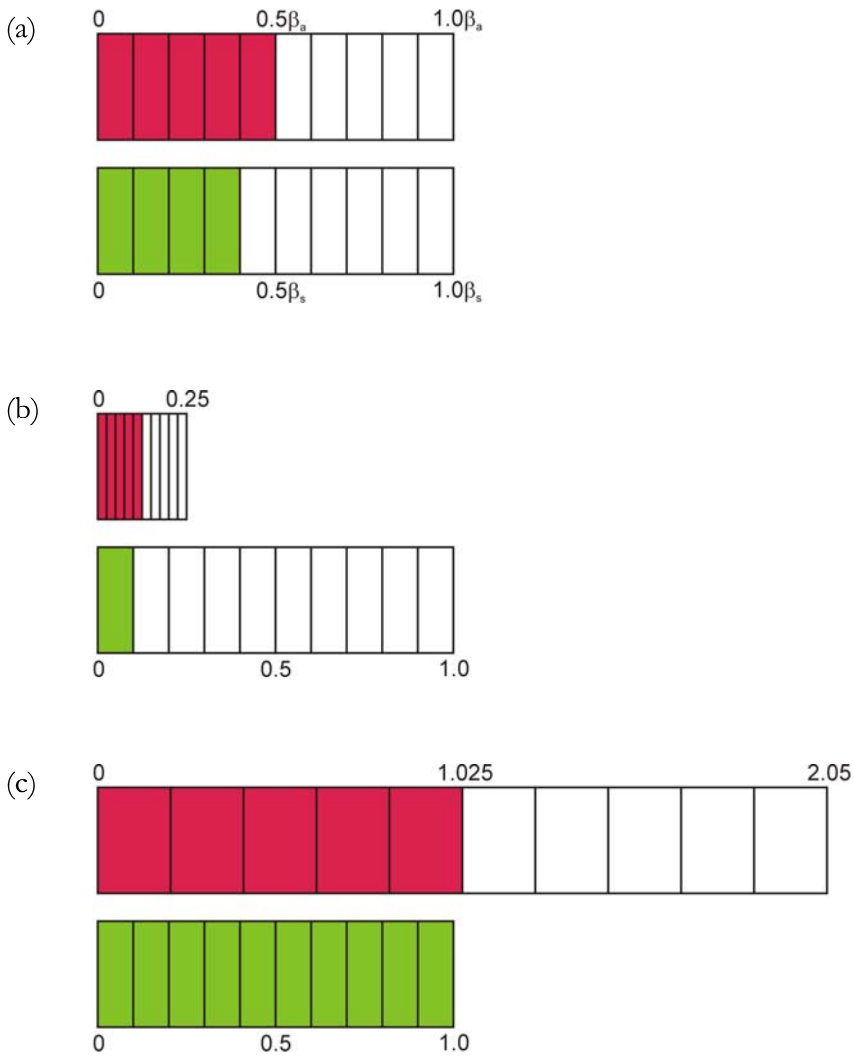

Periods

The number of periods over which to run the model was chosen so that each repeat of the model was given ample time to reach the maximum contagion determined by the parameterisation. The program was set up so that warning flags would be recorded if the model did not reach equilibrium by the end of the steps.

Repeats

This was based on establishing standard deviation of cascades.

3 The Data

We examine evidence from four separate industries: the Financial & Professional Services; Health Sciences; Creative/Digital/New Media/IT/Communications and Engineering & Textiles industries in the Manchester City Region (MCR). These industries span quite different types of industrial activity and so provide a broad evidence base. In addition, using data from different industries enables us to examine contrasts between industries in the drivers of innovation. The primary research was conducted as part of a broader study to examine the prospects and policies of the Manchester City Region[11]. This information was used to calibrate the agent based model described in Section 2 for firms in each of the industries.

The firms are connected by a network, and direct flows of innovation can only take place between firms which are connected. The particular structure of the network is based on the evidence of the business relationships between firms.

Drawing on the primary research, we obtain the distributions across firms of the networks and the three key parameters of the innovation model. These are:

- The propensity of a firm to innovate
- The secrecy index, which is the degree to which firms seek to protect their innovation
- The absorptive index, which reflects the willingness of firms to discover and participate in innovation.

Given the distributions across firms of these three key parameters and the structure of the networks across which firms are connected, we are able to analyse the propensity of innovation to spread across the firms of an industry, and to identify the parameters to which the spread is most sensitive.

3.1 The Network Structure

[11] The Manchester Independent Economic Review is a group of several projects looking at various aspects of the City's economy. The models described here have been developed in one of these projects.

The network structures used in the model for each industry are derived from the relations between firms in the MCR identified from the surveys. Organisations were asked what proportion of their purchases and their sales are made within the Manchester City Region. The responses for both questions had similar distributions so they were combined into a single entity.

The resulting distribution for each industry gives us the proportion of firms that perform a specific percentage of their business within the MCR, both upstream and downstream in their supply chain, but not what absolute level this represents.

These distributions are used to calibrate the relative degree of connectivity or links between firms in the network model for each industry. Note that this implies that organisational size does not affect the degree of connection with other organisation. Evidence suggests that larger firms focus their supply chains so that fewer firms are involved, while developing stronger relationships with those that remain, which would be consistent with this[12].

Supply chain relationships are important to the spread of innovation between firms. Dyer and Singh (1998) argue that "a firm's alliance partners are, in many cases, the most important source of new ideas and information that result in performance-enhancing technology and innovation"[13].

Table 3.1 shows the mean percentage of business performed within the MCR for each of the four industries while figures 3.1-3.4 show the overall distributions.

Figures 3.1-3.4 demonstrate that a surprisingly high proportion of firms identified themselves as performing no proportion of their business within the MCR. In the network model this is equivalent to an organisation having no links. Since the aim of the network study is to understand the drivers that enable the innovation of one firm to

---

[12] Rhodes, E., Warren, J., & Carter, R., (2006), *Supply Chains and Total Product Systems: A Reader*, Blackwell Publishing

[13] Dyer, J H. & Singh, H. (1998), The relational view: cooperative strategy and sources of inter-organizational competitive advantage, *Academy of Management Review*

spread to others, these isolated firms are unable to spread or receive an innovation and are therefore redundant.

Table 3.1: Mean percentage of links within the MCR for organisations within each industry

Industry	Mean % of business links within the MCR	
	Including those that do not have any MCR links	Excluding those that do not have any MCR links
Creative/Digital/New Media/IT/Communications	28%	43%
Financial & Professional Services	38%	52%
Health Sciences	53%	70%
Engineering & Textiles	22%	34%

These distributions do not allow us to know how many links a given firm within the MCR actually has; it simply tells us the relative connectivity of the population of firms. For example, in the case of the Health Sciences industry we know that there are three times as many firms with all of their business relationships contained within the MCR as there are firms who have 26-50% of their links within the region. We do not know how many links this corresponds to.

Note also that the firms' networks do not distinguish here between the kinds of linkages – they are the linkages of firms from a particular kind of business but with all kinds of organisations.

Comparison Between Questionnaire Derived and Model Degree Distributions

In practice the distributions derived from the questionnaire are not perfectly replicated using the *k*-clique algorithm discussed earlier. Residuals may exist in the distribution that cannot form a perfect *k*-clique. Table 3.2 shows typical values for the percentage of links that this method successfully incorporated into the network as a proportion of the total in the questionnaire derived distribution; this is called the degree of reciprocation.

Table 3.2 indicates that the algorithm is capable of assigning over 90 per cent of the links in all but the densest networks (where S_0 is 64); even then it is above 85 per cent. As S_0

decreases and the network structure becomes less dense the algorithm improves in efficiency.

There are two options at this point regarding how to deal with the residual differences between the questionnaire derived and model degree distributions. First, since the degree of distribution is adequately high across all parameterisations the differences could be ignored, second, the difference could be removed by including un-reciprocated links (links where an innovation can only be adopted in one direction).

The second method was chosen so that the degree distribution in the model exactly matched that of the questionnaire but as a result a residual proportion of these were unreciprocated.

Table 3.2: Analysis of the performance of the network generating algorithm for each industry distribution and value of S_0.

Industry	S_0	Typical Degree of Reciprocation (Target of 100%)	Average Randomisation (Target of 50%)	Average Super-Component Size
Creative/Digital/New Media/IT/Communications	8	99.2%	49.96%	82.2%
	16	98.5%	49.97%	100.0%
	24	97.7%	-	100.0%
	32	95.6%	49.97%	100.0%
	64	91.7%	-	100.0%
Financial & Professional Services	8	98.5%	49.93%	94.5%
	16	98.2%	49.97%	100.0%
	24	97.9%	-	100.0%
	32	97.5%	49.97%	100.0%
	64	88.6%	49.96%	100.0%
Health Sciences	8	99.8%	49.96%	97.5%
	16	98.3%	49.96%	100.0%
	24	97.5%	-	100.0%
	32	96.1%	-	100.0%
	64	97.0%	-	100.0%
Engineering & Textiles	8	97.9%	49.94%	67.8%
	16	96.1%	49.98%	100.0%
	24	96.0%	-	100.0%
	32	96.9%	49.97%	100.0%

	64	86.4%	-	100.0%

Network Super-Component Size

One consideration with the method developed to generate the network structures is that there is a probability that not all of the 1000 firms in the model end up connected into a single component. This increases in likelihood as S_0 is reduced. It is important to understand what the average maximum component size is for each parameterisation of the network.

For example if on average when S_0 is 8 the largest connected component, or super-component, is 80 per cent of all of the agents in the model, then this acts as a ceiling to the maximum percolation that a single innovation can achieve.

The average super-component size was calculated for each of the five network parameterisations used for each industry. The results of this analysis is shown in Table 3.2, it shows that the super-component size only needs to be considered when the industry networks are at their lowest density or when S_0 is 8. The relative size of the industry super-components under this parameterisation is simply a restatement of the average number of links presented in Table 3.1.

Figure3 1: Distribution of links in the Creative/Digital/New Media/ICT Digital/Communications Industry

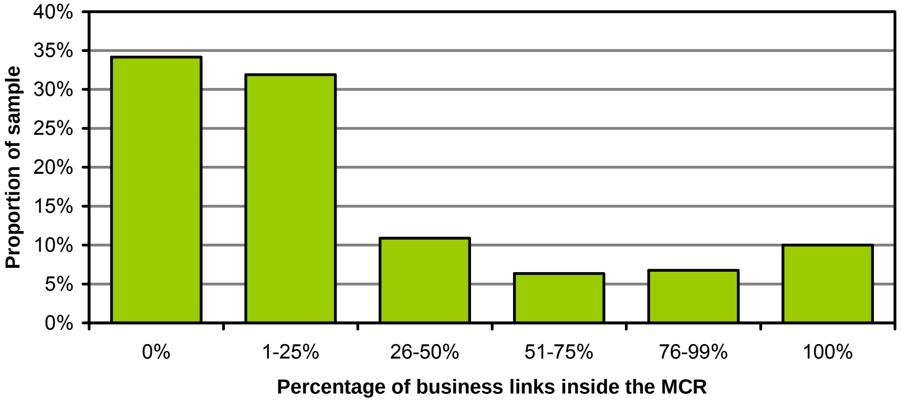

Figure 3.2: Distribution of links in the Financial & Professional Services Industry

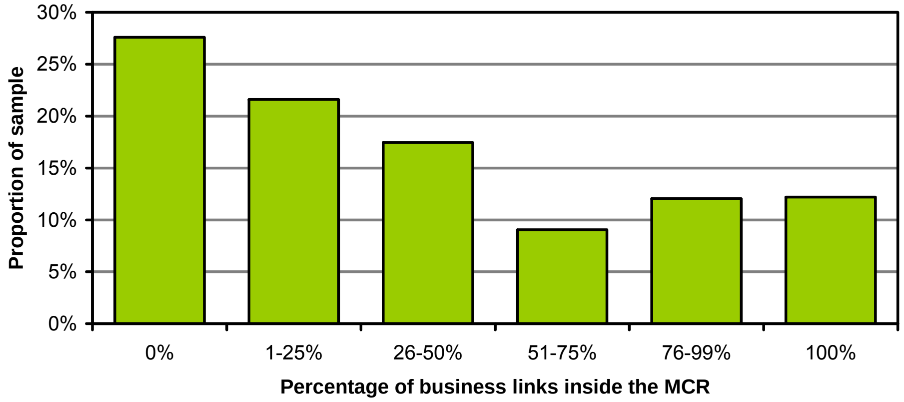

Figure 3.3: Distribution of links in the Health Sciences Industry

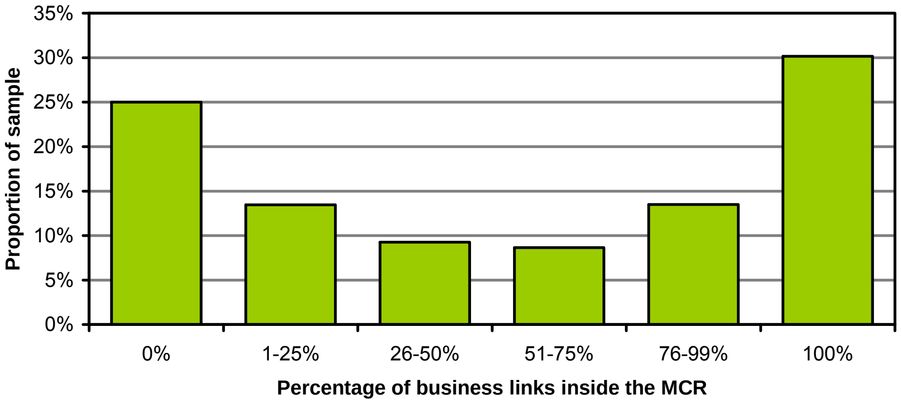

Figure 3.4: Distribution of links in the Engineering & Textiles Industry

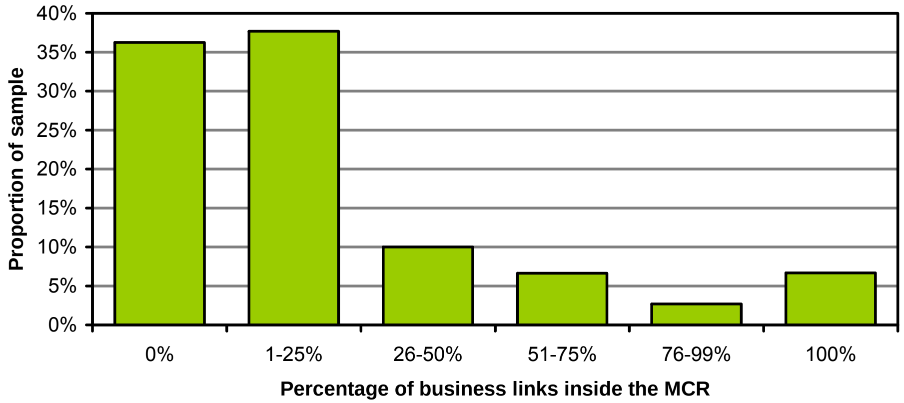

3.2 Agent Characteristics

The Secrecy Index or Innovation Protection

The secrecy index in the model measures the strengths of the barriers to the spread of innovation. The market pressures that lead firms to innovate will also mean they attempt to prevent the spread of the innovation to competitors. The literature shows the importance of the structures of knowledge governance to the spread of innovation[14]. If firms protect innovations, the ability of a network structure to spread innovation is reduced.

In order to construct the distribution for the secrecy index of organisations in the model the survey asked firms about the importance they placed on various methods for protecting innovations. These protective measures were:

- Registration of design
- Trademarks
- Patents
- Confidentiality agreements
- Copyright
- Secrecy
- Complexity of design
- Lead time advantage on competitors

From these answers we can build a picture of how active firms in an industry are in protecting their innovations, and thus part of the picture of how easy it is for innovations to spread. The exact methodology is presented in Appendix I. Figures 5-8 present the resulting distributions of the secrecy indices for each industry.

Table 3.3 presents the results of Kolmogorov-Smirnov tests which evaluate the statistical probability that the derived distributions are samples from the same underlying

[14] Antonelli, C. (2008) *Localised Technological Change: Towards the Economics of Complexity*, Routledge

distribution. The Health Sciences industry is clearly distinct to the others, but there is a 73 per cent chance that the Creative/Digital/New Media/IT/Communications and Financial & Professional Services' distributions are in fact the same.

Table 3.3: p-values of Kolmogorov-Smirnov tests to evaluate whether any two secrecy index distributions are statistically equivalent. Those cells shaded red indicate distributions that are statistically dissimilar at the level shown, those un-shaded cells are statistically similar at the level shown.

	Creative/Digital/New Media/IT/Communications	Financial & Professional Services	Health Sciences	Engineering & Textiles
Creative/Digital/New Media/IT/Communications	-	0.37	<0.01	0.74
Financial & Professional Services	0.37	-	0.01	0.74
Health Sciences	<0.01	0.01	-	<0.01
Engineering & Textiles	0.74	0.74	<0.01	-

Table 3.4 gives an average secrecy index of each industry, a high secrecy score indicates an industry in which there are higher barriers to the sharing of innovations, and a lower score means there are lower barriers to sharing.

Table 3.4: Mean secrecy index of each industry

Industry	Mean Secrecy Index
Creative/Digital/New Media/IT/Communications	0.34
Financial & Professional Services	0.31
Health Sciences	0.20
Engineering & Textiles	0.34

Figure 3.5: Distribution of secrecy index in Creative/Digital/New Media/ICT Digital/Communications

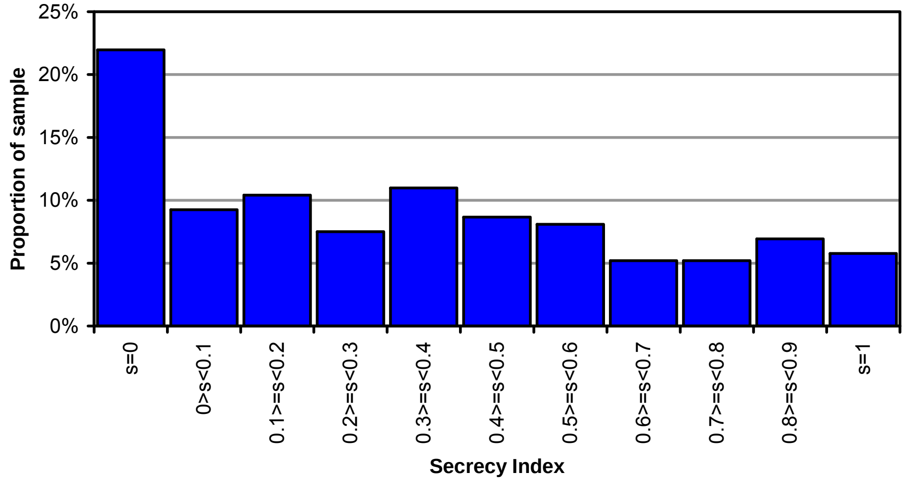

Figure 3.6: Distribution of secrecy index in Financial & Professional Services

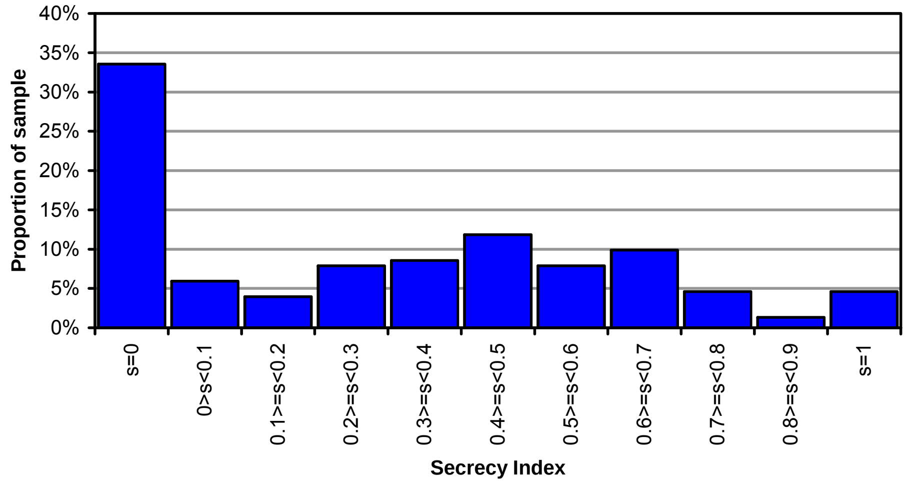

Figure 3.7: Distribution of secrecy index in Health Sciences

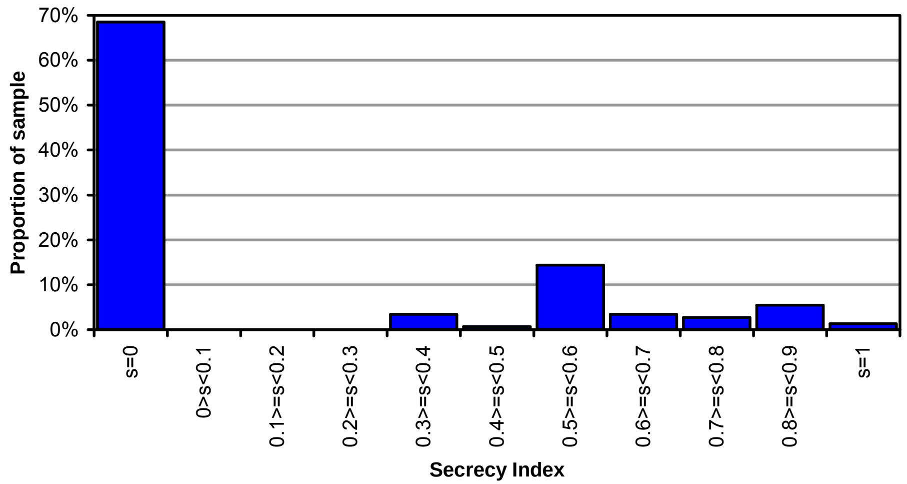

Figure 3.8: Distribution of secrecy index in Engineering & Textiles

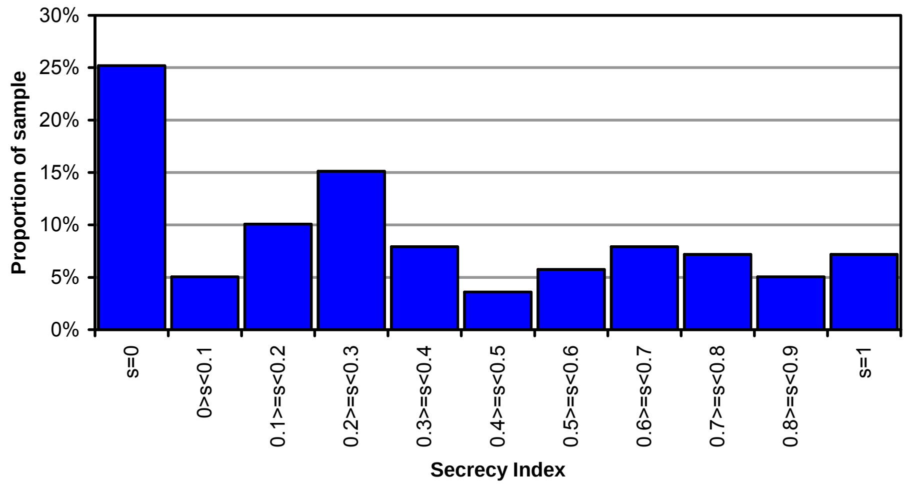

Absorptive Index or Pursuing Innovation

Spreading an innovation using the exchange process is a two way relationship – the firm with the innovation must be willing to share it, and the firm without the innovation must be actively pursuing and able to adopt it. The ability to absorb an innovation depends on a number of factors:

- The knowledge parity between the firm with and without the innovation
- The importance the firm without the innovation places on absorbing innovations
- The involvement of a firm in wider networks – knowledge transfer partnerships, business networks, links with higher education institutions, trade organisations, cluster organisations, and so on.

The survey asked firms about their involvement in the final factor, as this is the most relevant from a network perspective, and from this we are able to construct an absorptive index for each firm. The literature showed that the ability to absorb innovation was intimately tied to firms' involvement in the network of firms around them[15]. The absorptive methods asked about were:

- Publicly-funded joint research programmes
- Other research associations
- Joint R&D agreements
- Licensing and second sourcing agreements
- Knowledge Transfer Partnerships and hosting research studentships
- Other links with Higher Education
- Member of/attends: trade associations, cluster organisations, business networks
- Member of other general business organisations

The exact methodology for calculating the index for an individual firm is detailed in Appendix II. Figures 3.9-3.13 present the resulting distributions of the index for each industry.

[15] For example: Cowan, R. (2004) Network models of innovation and knowledge diffusion, *Merit* and *Infonomics*

Table 3.5 again presents the results of Kolmogorov-Smirnov tests which evaluate the statistical probability that the derived distributions are samples from the same underlying distribution. There is a 73 per cent chance that the Health Science industry's distribution and both the Financial & Professional Services and the Creative/Digital/New Media/IT/Communications' distributions are in fact the same.

Table 3.5: p-values of Kolmogorov-Smirnov tests to evaluate whether any two absorptive index distributions are statistically equivalent. Distributions are statistically similar at the level shown.

	Creative/Digital/New Media/IT/Communications	Financial & Professional Services	Health Sciences	Engineering & Textiles
Creative/Digital/New Media/IT/Communications	-	0.74	0.37	0.74
Financial & Professional Services	0.74	-	0.37	0.99
Health Sciences	0.37	0.37	-	0.99
Engineering & Textiles	0.74	0.99	0.99	-

As for the secrecy index we can give a mean absorptive index for each industry to give an initial description of the ability of firms within that industry to absorb innovations, this is shown in Table 3.6.

An absorptive capacity of one indicates a firm that undertakes every method of absorbing innovations that the survey identified. Such a firm has a very high ability is absorb innovations from other firms it is connected with. An absorptive capacity of zero indicates a firm that undertakes none of the methods of absorbing innovations. Such a firm is unable to extract innovations from the wider network of firms in which it is embedded.

Table 3.6: Mean absorptive index of each industry

Industry	Absorptive Index
Creative/Digital/New Media/IT/Communications	0.18
Financial & Professional Services	0.22
Health Sciences	0.14
Engineering & Textiles	0.17

Propensity to Innovate

The final distribution derived from the survey was the propensity of organisations to innovate in each industry. Firms were asked whether or not they had engaged in a range of activities conducive to creating an innovation, over the three year period 2005-2007. Specifically, they were asked whether they had performed the following:

- In-house R&D
- Acquisition of external R&D
- Acquisition of machinery, equipment and software
- Acquisition of external knowledge
- Training to support new products, services, or business and process improvement
- Design expenditure
- Marketing and advertising specifically related to new and improved services or products
- Major changes in business structure, practices or processes

Figure 3.9: Distribution of absorptive index in Creative/Digital/New Media/ICT Digital/Communications

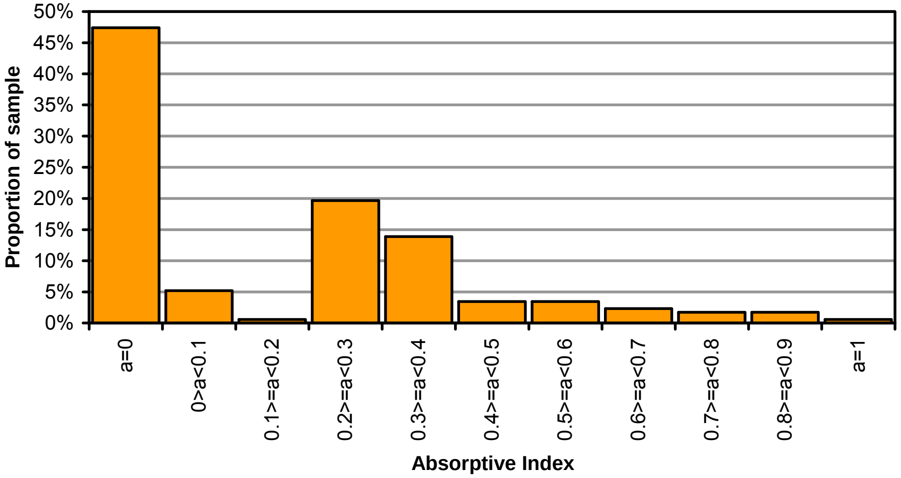

Figure 3.10: Distribution of absorptive index in Financial & Professional Services

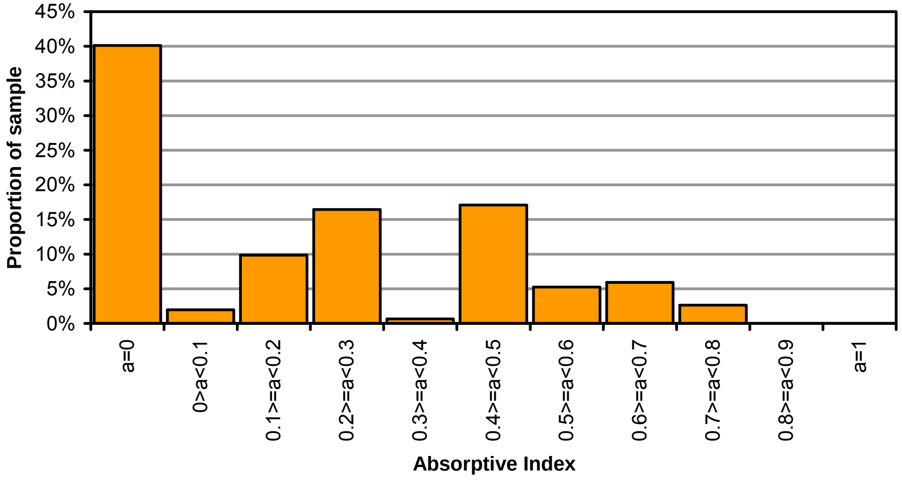

Figure 3.11: Distribution of absorptive index in Health Sciences

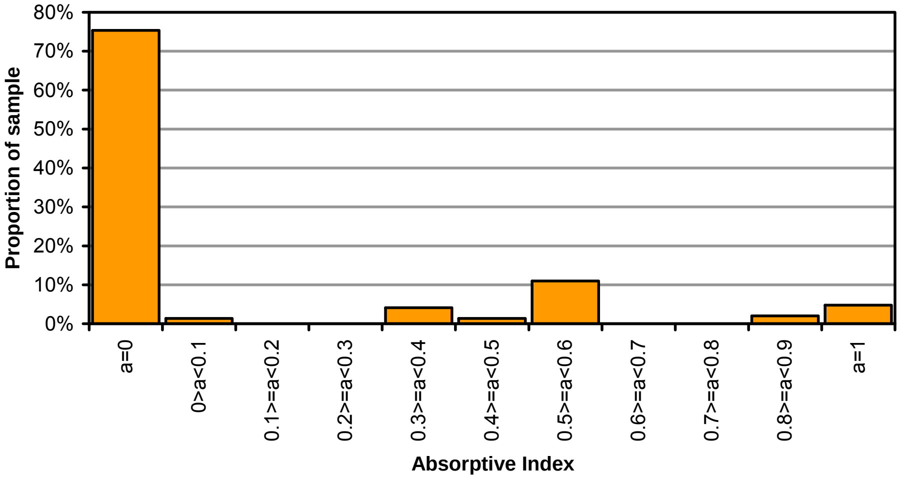

Figure 3.12: Distribution of absorptive index in Engineering & Textiles

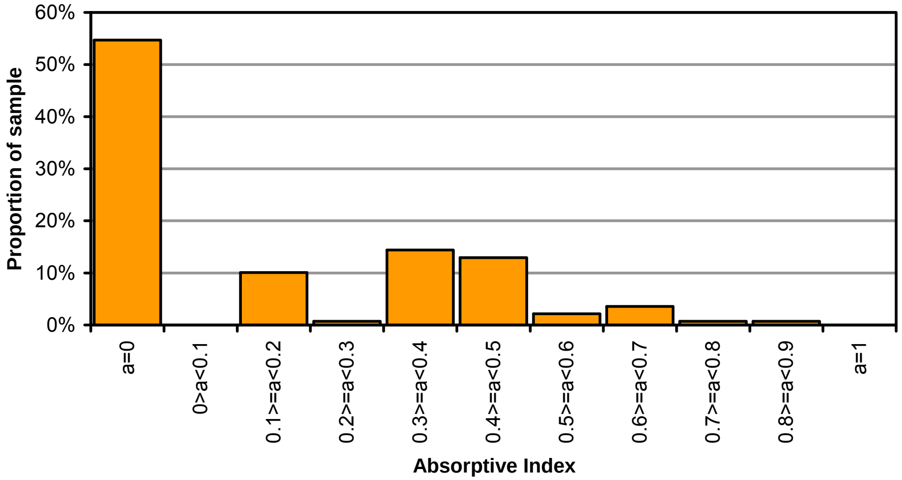

From the responses a propensity to innovate was constructed for each firm. The method of calculation is detailed in the Appendix.

Table 3.7 presents the results of Kolmogorov-Smirnov tests which evaluate the statistical probability that the derived distributions are samples from the same underlying distribution. The Health Sciences industry is clearly distinct from the Creative/Digital/New Media/IT/Communications and Financial & Professional Services' distributions but there is an 85 per cent probability that it is the same as the Engineering & Textiles distribution.

Table 3.7: p-values of Kolmogorov-Smirnov tests to evaluate whether any two innovation index distributions are statistically equivalent. Those cells shaded red indicate distributions that are statistically dissimilar at the level shown, those un-shaded cells are statistically similar at the level shown.

	Creative/Digital/New Media/IT/Communications	Financial & Professional Services	Health Sciences	Engineering & Textiles
Creative/Digital/New Media/IT/Communications	-	0.99	0.05	0.74
Financial & Professional Services	0.99	-	0.05	0.74
Health Sciences	0.05	0.05	-	0.15
Engineering & Textiles	0.74	0.74	0.15	-

An innovation index close to one indicates a firm that engages in a high number of innovation-generating activities. An index close to zero indicates a firm that undertakes very few innovation related activities. Table 3.8 shows the mean innovation index of each industry.

Table 3.8: Mean innovation index of each industry

Industry	Innovation Index
Creative/Digital/New Media/IT/Communications	0.50
Financial & Professional Services	0.48
Health Sciences	0.29
Engineering & Textiles	0.51

Figures 3.13-3.16 show the distribution of innovation index of each firm within an industry. This is only relevant in the second of the two network models when more than one innovation can be generated and spread amongst firms.

Figure 3.13: Distribution of innovation index in Creative/Digital/New Media/ICT Digital/Communications

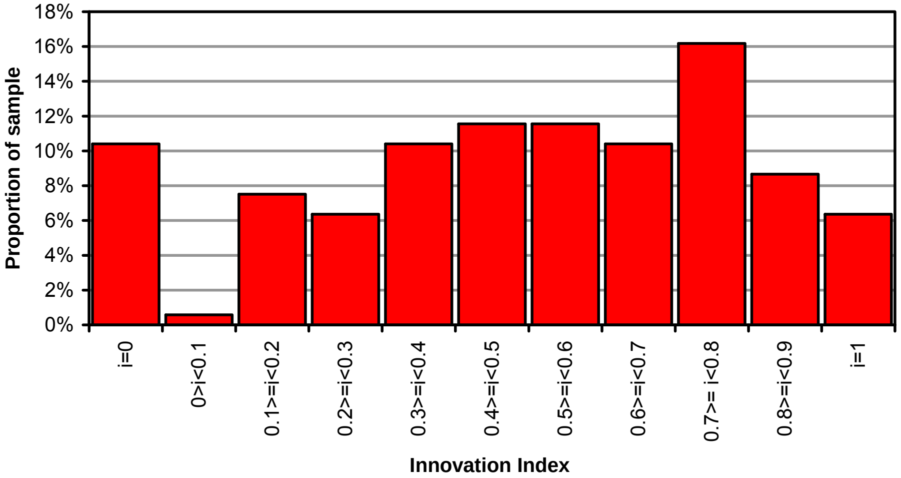

Figure 3.14: Distribution of innovation index in Financial & Professional Services

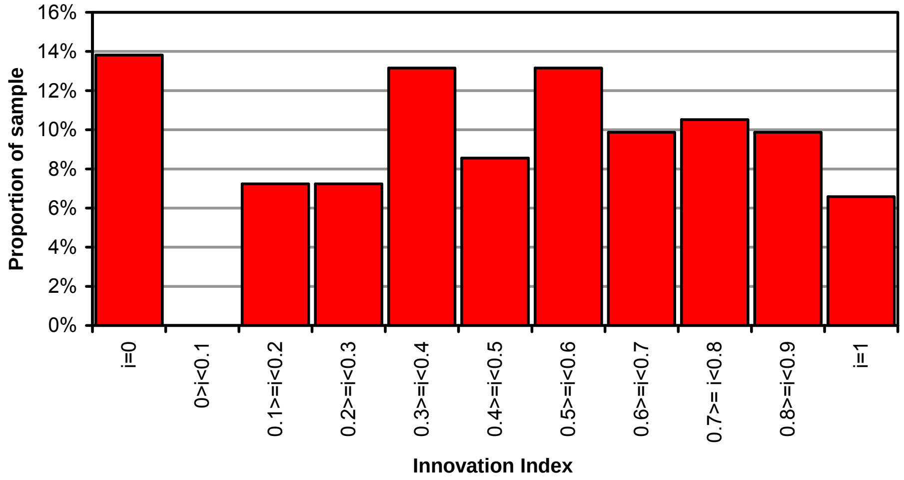

Figure 3.15: Distribution of innovation index in Health Sciences

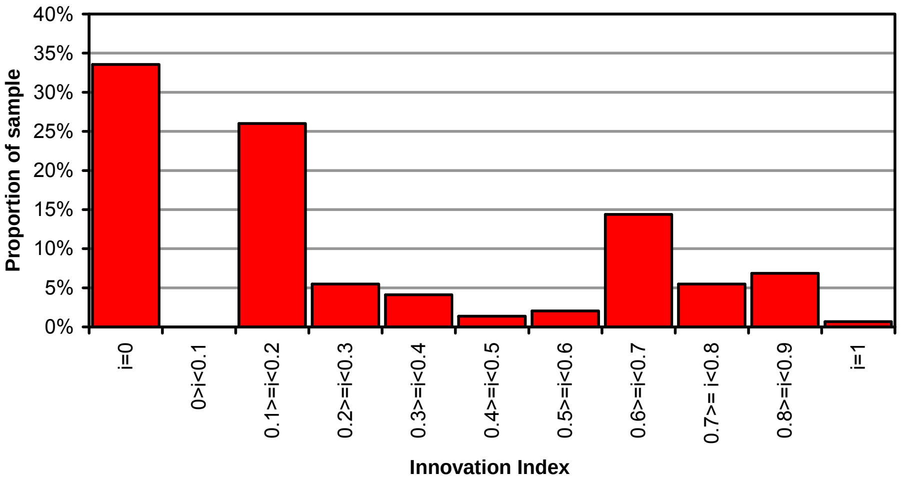

Figure 3.16: Distribution of innovation index in Engineering & Textiles

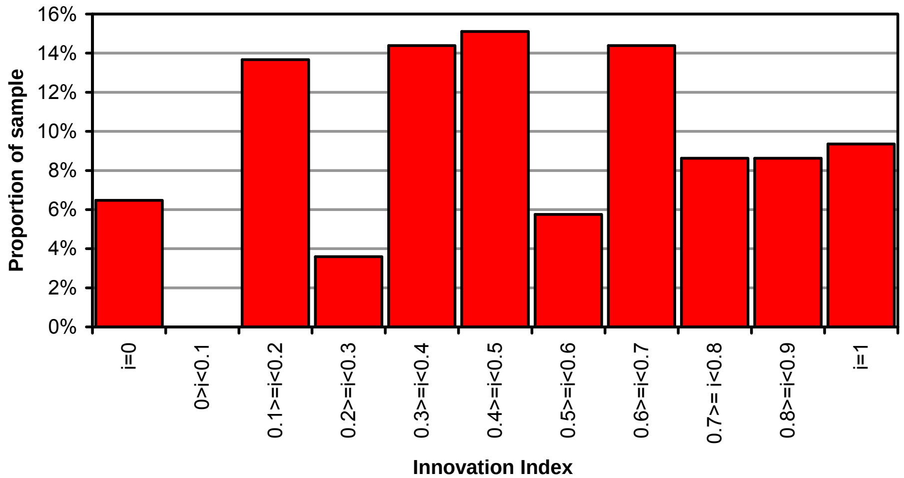

4 Results

4.1 Cascades

The output of a single execution of the model is a time series of the percentage of firms that adopted the original innovation. Figure 4.1 shows two examples for the Creative/Digital/New Media/IT/Communications industry using the same parameterisation. It shows two contrasting examples: first, the green line where over 50 time steps the innovation spread to 3.5 per cent of the firms in the network and second, the red line where a much larger contagion occurred and 75 per cent of the agents adopted the innovation over the same number of steps.

Figure 4.1: Example results for single runs of the model for the Creative/Digital/New Media/IT/Communications industry (α=0, ε=1, S_0=16, β_a=0.25, β_s=1)

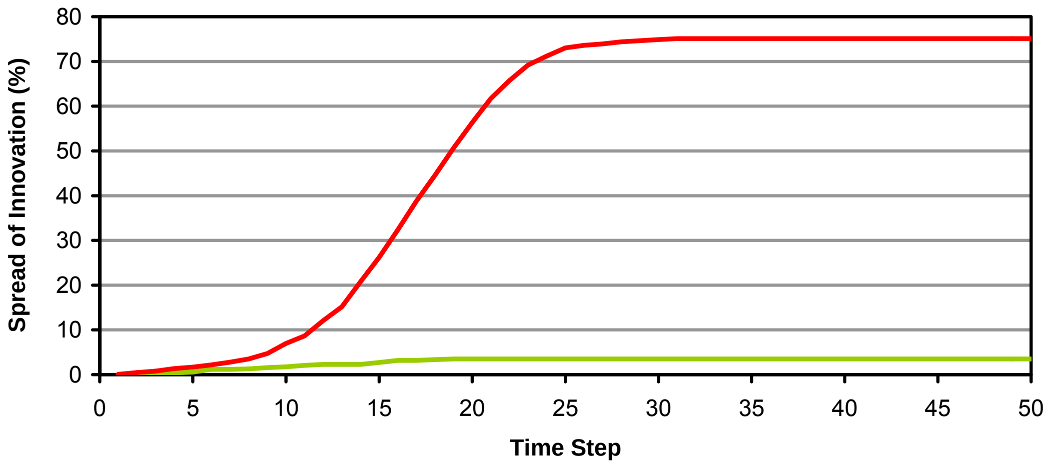

Figure 4.1 demonstrates that even for the same parameterisation the results of individual runs of the model can be very different. This variation is a result of differences in the arrangement of the network such as the degree of connection of the firm that generates the innovation as well as the randomly assigned personal thresholds, absorptive and secrecy indices. It is unrealistic to expect to be able to map the real world network and assign each firm their actual behavioural characteristics so this process of randomisation is necessary to explore the possible impacts of the myriad of arrangements. It is for this reason that the model is repeated 1000 times for each parameterisation, for each industry. The final spread of the innovation (3.5 and 75 per cent in the examples shown in Figure

4.1) across these 1000 repeats is plotted as a histogram for each industry parameterisation, as shown in Figure 4.2.

Figure 4.2: Histogram of the final spread of innovation from 1000 runs of the model for the Creative/Digital/New Media/IT/Communications industry (α=0, ε=1, S_0=16, β_a=0.25, β_s=1)

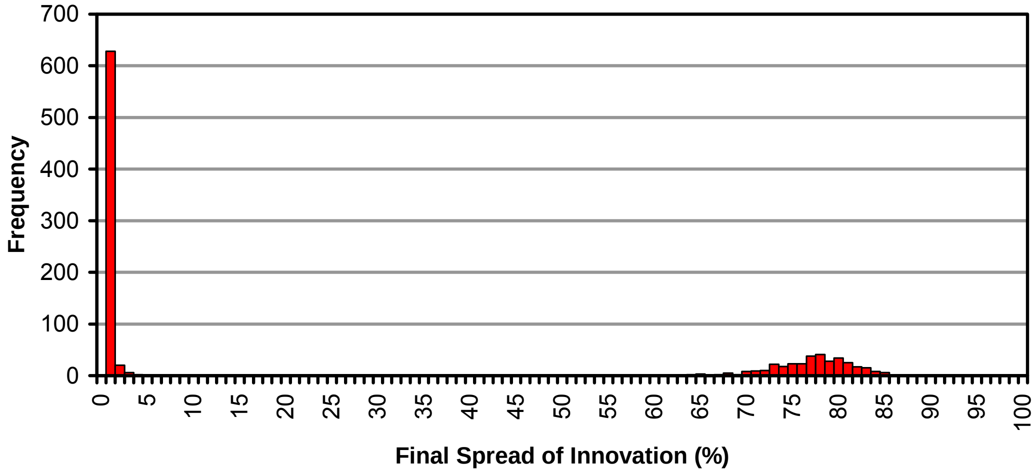

Figure 4.2 shows that the cascades that occur form two distinct distributions, it is the upper symmetric distribution that is of interest for this study which we will call the global cascade (as opposed to the lower skewed distribution which we call the local cascade). As discussed the aim of this analysis is to understand the drivers that create this wider spread contagion, the assumption being that the Manchester City Region will benefit if the innovations of one firm are enabled to spread to others in the regional economy.

In order to evaluate these drivers, the characteristics of the upper symmetric distribution need to be captured, these characteristics include its mean, standard deviation and proportion of the 1000 repeats it includes. This analysis focuses on the first of these. In order to evaluate the mean of the distribution the lower skewed distribution needs to be removed from the data.

4.2 Significance of Global Cascade Mean

Before moving onto into the more substantive results it is important to evaluate the statistical robustness of the mean of the global cascade. For example the mean value would vary significantly when comparing runs of the model that only used ten repeats in the averaging. If we repeat the model 1000 times how different could we expect the results of this be compared to a second run of 1000 repeats?

Figure 4.3 shows an analysis of the standard deviation of the mean of the global cascade across ten runs of the model, for different number of repeats. This standard deviation is expressed as a percentage of the mean. It shows that the standard deviation decays as a power law. When using 1000 repeats, 68 per cent of the time the mean can be expected to vary by less than 0.1 per cent[16]. Of course this only applies to this parameterisation but even if the variation were an order of magnitude higher, this level of accuracy would be acceptable.

Figure 4.3: Standard deviation of the mean of ten runs of the model using different numbers of repeats for the Creative/Digital/New Media/IT/Communications industry (α=0, ε=1, S_0=16, β_a=0.25, β_s=1)

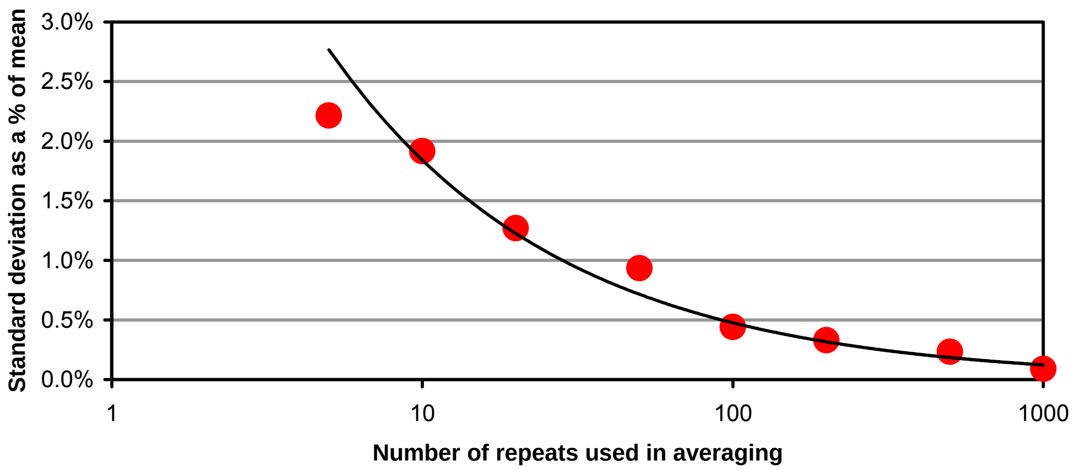

[16] And only 0.2 per cent 95 per cent of the time, assuming that the mean of the global cascade distribution follows a normal distribution.

Table 3.9: Comparison of the mean global cascade, its standard deviation and the proportions of runs of the model that result in global cascades for the four industries using a single parameterisation (α=0, ε=1, S_0=16, β_a=0.25, β_s=1).

Industry	Mean Global Cascade (%)	Standard Deviation (%)	Proportion of runs that are global cascades (%)
Creative/Digital/New Media/IT/Communications	76.7	4.2	34.3
Financial & Professional Services	92.8	1.8	55.2
Health Sciences	81.3	4.5	71.4
Engineering & Textiles	57.0	8.0	22.6

Table 3.9 illustrates the mean global cascade and timeframe characteristics of the model across the sectors for this particular set of parameters. It highlights that the different industry survey distributions have a significant impact on the characteristics of innovation contagion. It is probably most helpful to think of the sectors as capturing variation of performance of firms of different types as much as different industries

4.3 Network Descriptions

Betweeness and Average Shortest Path Length

There is an extensive range of characteristics associated with the structure of a network; two relevant ones for the innovation model are shortest path length and betweeness. The first is a network parameter called the average geodesic distance which measures the shortest path between each pair of firms, in terms of number of links, and finds the average. Second is betweeness, a parameter associated with each agent that measures the frequency of involvement of that firm in the shortest path between all other pairs of firms. This is a therefore a measure of the presence of key actors in a network a higher maximum betweeness indicates the presence of a key actor that dominates the shortest paths between firms.

Summary results for these characteristics are shown in Table 4.1, it shows that as the density of links increases in a network both the maximum betweeness and the average shortest path length decrease. The Health Sciences industry is identified as the network structure that has both the shortest average path length and lowest maximum in betweeness. On this measure then this industry appears to have a network structure conducive to spreading an innovation,

Table 4.1: Analysis of the maximum betweeness and average shortest path length for each industry distribution and value of S_0.

Industry	S_0	Maximum Firm Betweeness	Average Shortest Path Length
Creative/Digital/New Media/IT/Communications	8	25247	5.7
	24	12610	3.5
	64	7290	2.6
Financial & Professional Services	8	15423	5.3
	24	7273	3.2
	64	5072	2.4
Health Sciences	8	11615	4.5
	24	4827	2.9
	64	3100	2.2
Engineering & Textiles	8	45608	7.2
	24	17923	3.8
	64	11653	2.7

4.4 Results for the Exchange Mechanism

Figure 4.4 shows how the mean of the global cascade varies for each industry when the only method by which an innovation can be adopted is through the exchange mechanism.

- The model captures the impact that differences in both structure and behaviour between the industries on the mean size of a global cascade.

- Increasing the density of links (S_0) increases the spread of an innovation. The strength of this increase varies between industries with Health Sciences Financial & Professional Services seeing less sensitivity to the density of the network than the others.

- The most significant driver of producing larger global cascades is the absorptive capacity scalar. As this increases the cascades become larger, although the effect is minimal in the Health Sciences industry. This effect is strongest in the Creative/Digital/New Media/IT/Communications and Engineering & Textiles industries.

- The Health Sciences industry is the least responsive to changes in the absorptive capacity scalar. This is for a range of reasons:

    - The industry has the highest proportion of organisations with absorptive capacities of zero (around 75%). These firms are unaffected by increases in β_a.
    - In comparison, the industry network has a disproportionate number of organisations in the category with the highest number of links. This is what drives the low average shortest path and betweeness. In this respect the network can be viewed as more cohesive where even at the lowest scaling factor there are already multiple routes between organisations for the innovation to spread. Increasing this scaling then has a proportionately smaller impact compared to the other industries.
    - Although the Health Sciences industry is characterised by a significant lack of absorptive capacity, sufficiently dense network links exists that the remaining 25% of firms share multiple paths through which innovations can reliably spread. This reliability is demonstrated by Table 3.9 which shows that the industry has the highest proportion of runs of the model that result in global cascades.

- The system is the most sensitive to changes in the absorptive capacity scalar when it is between 0 and 1. If we ignore the Health Sciences industry, which as we have said is less affected by changes in the scalar, the strength of the response for the other three can be understood in terms of the proportion of links in a network that can be used to spread an innovation (i.e. where the absorptive capacity at one end is greater than the secrecy index at the other).

    Table 4.3 shows that the Financial & Professional Services industry has 22% of its links that meet this criteria when β_a=0.25, this rises to 30% when β_a=0.75. This increase is less than the doubling shown by both the Engineering & Textiles and the Creative/Digital/New Media/IT/Communications (Table 4.3) and explains why their global cascades shown in Figure 4.4 respond much more strongly to changes in the absorptive capacity scalar between 0 and 1.

A simple indicator of the response of an industry to changes in the absorptive capacity scalar in this range is the difference between the mean absorptive and secrecy index of an industry.

- Figure 4.4 shows that when the absorptive scalar is greater than one the network scaling factor becomes more influential in driving how far global cascades can reach. Each system appears to have a maximum limit, if we again exclude the Health Services industry which we have covered already, this limit can be readily understood.

- If we take the Financial & Professional Services industry this shows the best average ability to spread innovation for a given parameterisation. However it also shows the second highest proportion of firms with an absorptive capacity of zero. This is mitigated by its higher levels of connectivity. This argument is exactly the same as the one outlined for the Health Services sector above. What we see is an antagonism between the influence of the behavioural distributions and the structural ones. These can mitigate each other in the middle ground but when one side wins we get a system like the Health Services sector and when the other does we get results like the Financial Services sector.

Table 4.2: Analysis of the percentage of links in an industry's network available for transmitting innovations using the exchange method.

Industry	**Proportion of links in the network where the exchange of an innovation is possible for a given value of absorptive capacity scalar**		
	β_a=0.25	β_a=0.55	β_a=0.75
Creative/Digital/New Media/IT/Communications	13.0%	20.0%	24.8%
Financial & Professional Services	21.9%	26.1%	30.3%
Health Sciences	16.9%	17.2%	18.7%
Engineering & Textiles	12.5%	18.4%	23.3%

Figure 4.4: Spread of a single innovation for different network scaling factors (S_0) *for each industry as a function of the absorptive capacity scalar* (β_a) *when no mimicking behaviour is allowed* (α=1.1, ε=1.1)

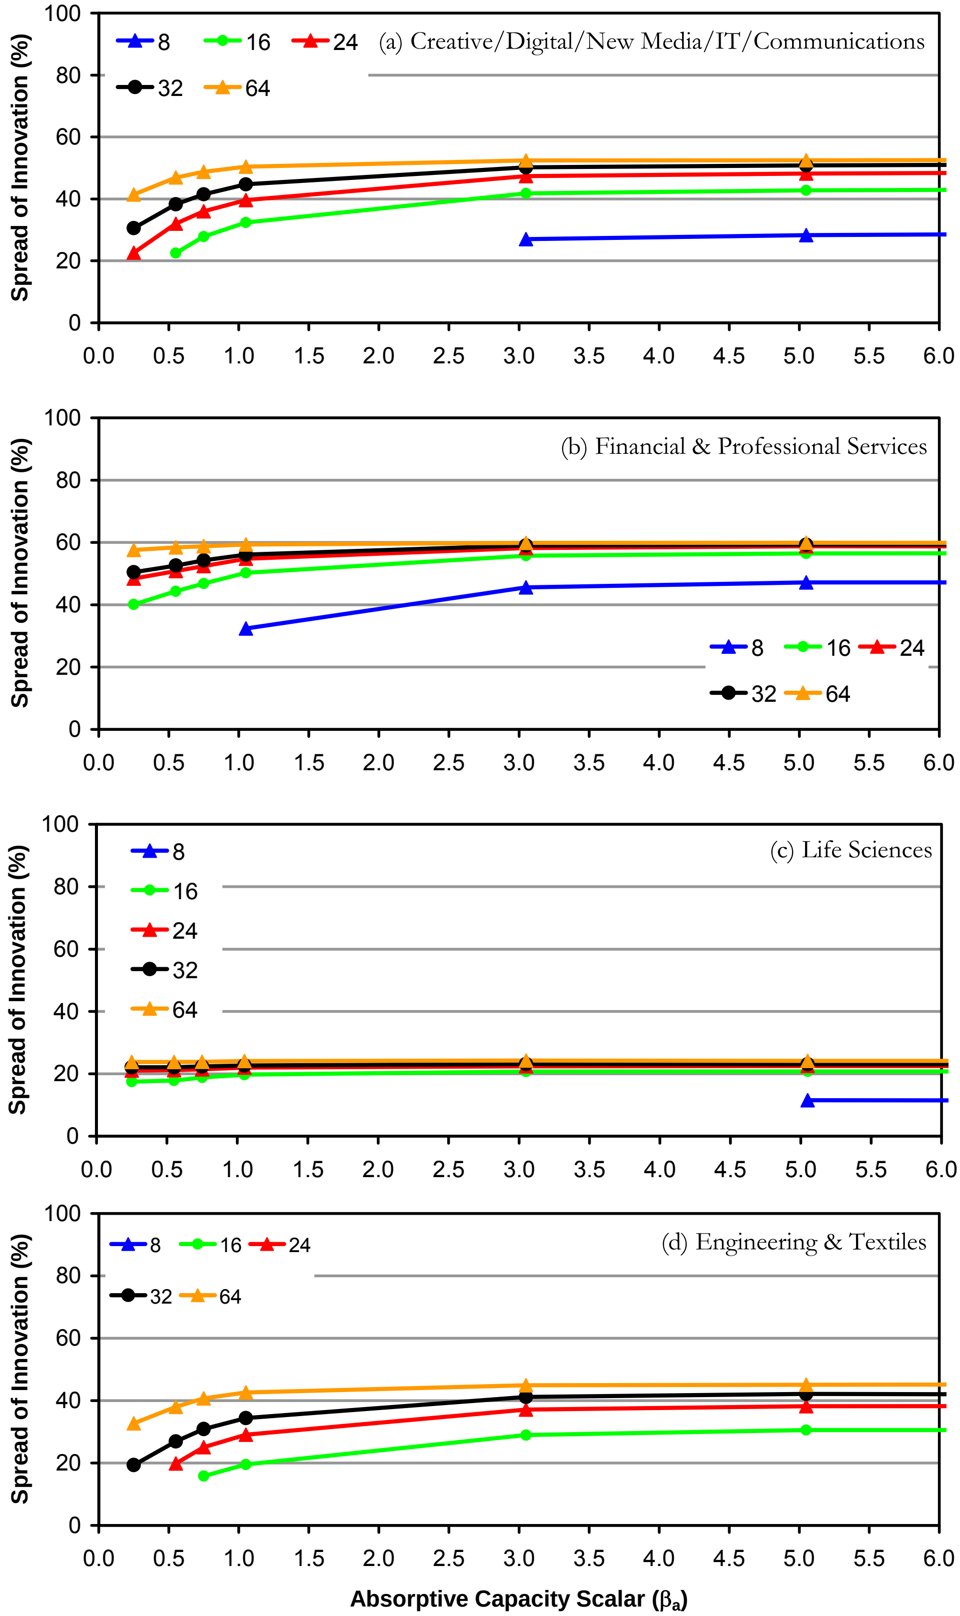

4.5 Spread of Innovation Through Copying Behaviour

Figure 4.5 shows an analogous set of results to Figure 4.4 but where innovation can only be adopted is through the mimicking mechanism. Note that lower α corresponds to greater ease of innovation.

- The copying behaviour appears incapable of creating the high cascades seen for the exchange behaviour.

- For all values of the minimum threshold the spread of an innovation is low and decays rapidly as it is increased.

- The conclusion is that copying behaviour is a poor method of spreading an innovation. In fact the relative strengths of the two ways in which innovation can spread were of the 100-100,000 to 1[17].

- The reason that copying behaviour is less successful is not because the parameterisations chosen in someway force this to be the case. When α is 0 one fifth of all firms in the model will only require less than a fifth of their links to possess an innovation in order for them to mimic their behaviour and adopt it themselves. This is comparable to the case when β_a=0.55 in Table 4.2. However, for the exchange relationship this leads to mean global cascades 10-30 times larger than for the copying behaviour. Again this demonstrates the relative effectiveness of the two mechanisms.

[17] In the case when either method could have spread the innovation in a given time step then this was counted in both the numerator and denominator in this calculation.

Figure 4.6: Spread of a single innovation for different network scaling factors (S_0) *for each industry as a function of the minimum threshold* (α) *when no exchange behaviour is allowed* (β_a=0)

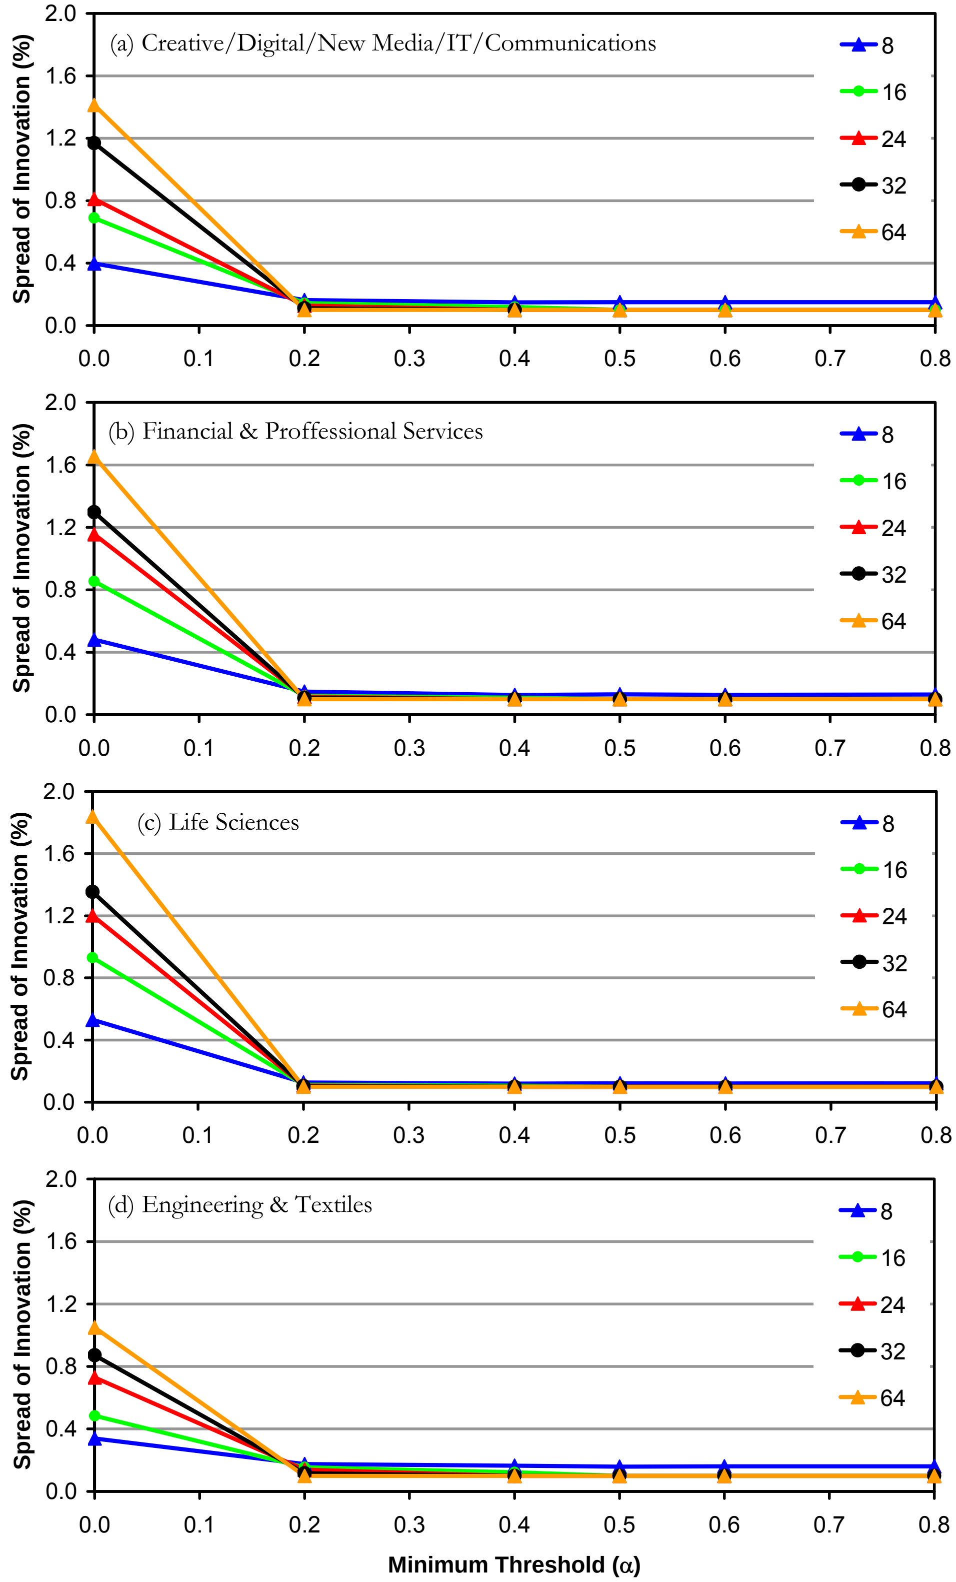

Spread of Innovation Through Both Behaviours: A Case Study

Figures 4.7 and 4.8 show the mean global cascade values when both the copying and exchange methods of adoption are allowed for different strengths of each. In figure 4.7 the parameterisation of the copying behaviour is held constant while the relative strength of the exchange behaviour is varied, while the opposite is true for figure 4.8. Only the Creative /Digital/New Media/ICT Digital/Communications industry is included in the interests of brevity although the conclusions are valid across all of the four industries.

- The exchange behaviour must be present for global cascades to occur.

- The exchange mechanism is not a replacement for the mimicking behaviour, together they achieve cascades that neither can achieve on their own. In some cases the addition of copying behaviour doubles the size of the cascade.

- Extremely large cascades can be achieved with even relatively low absorptive capacities, so long as there is strong copying behaviour. However, unless there is some absorptive capacity, copying will be insufficient.

- The real difference is that the addition of the copying behaviour removes some of the limits that we have discussed in Section 4.4. If an industry has a high proportion of firms with an absorptive capacity of zero then the copying behaviour allows for these organisations to still adopt an innovation *and* to subsequently pass it on. If an industry is relatively unconnected then the copying behaviour allows for multiple methods by which the innovation can spread amongst the existing links.

The net effect of adding the copying behaviour to the exchange based one is to increase the maximum possible sizes of the global cascades and reduce the sensitivity of the system to the absorptive capacity, this sensitivity was discussed in Section 4.4.

Figure 4.7: Spread of a single innovation for different network scaling factors (S_0) for the Creative /Digital/New Media/ICT Digital/Communications industry as a function of the minimum threshold (α) for various values of the absorptive scalar

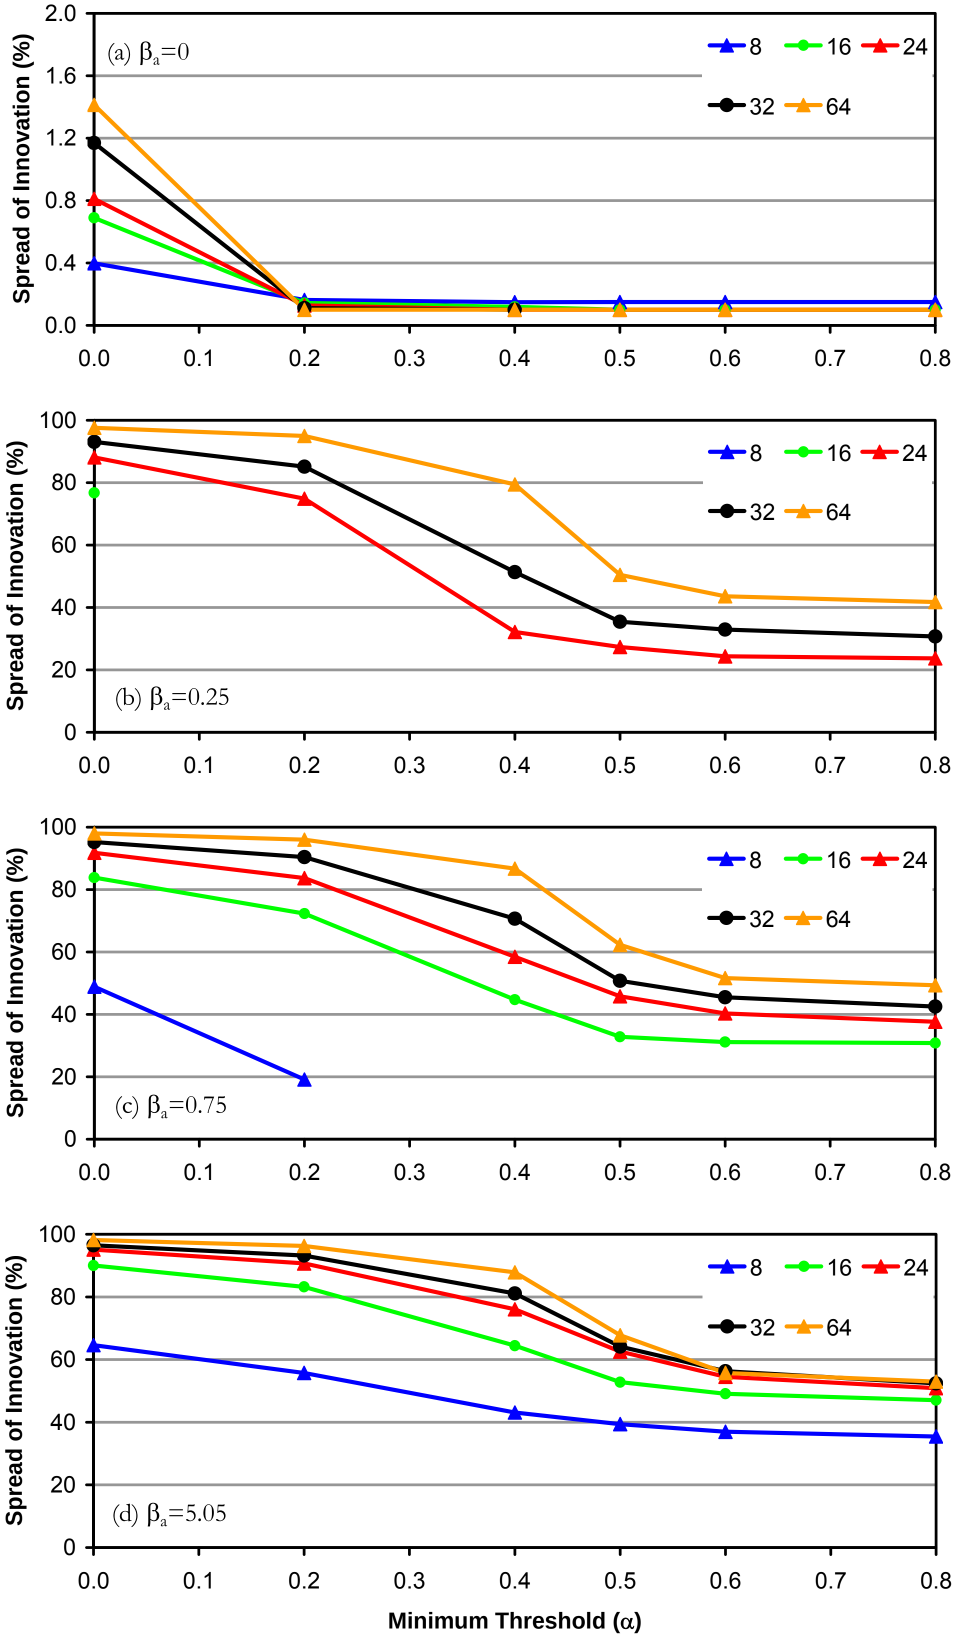

Figure 4.8: Spread of a single innovation for different network scaling factors (S_0) *for the Creative /Digital/New Media/ICT Digital/Communications industry as a function of the absorptive capacity scalar* (β_a) *for various values of the minimum threshold (a)*

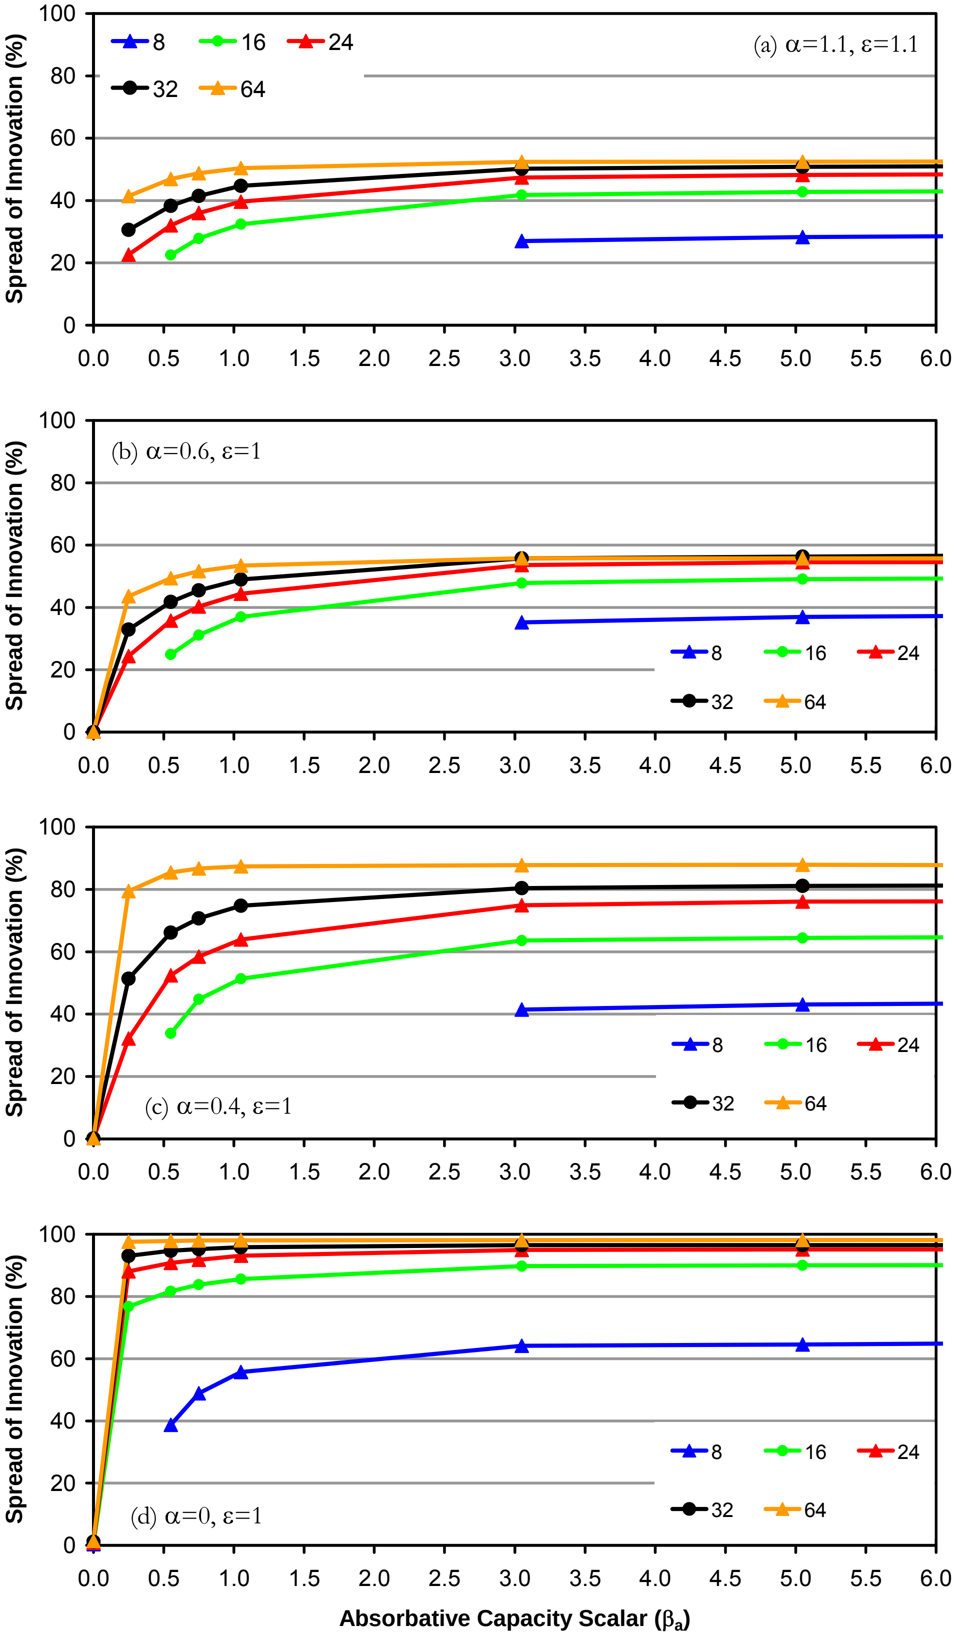

4.6 Length of Time to Spread Innovation

So far we have examined how the spread of innovation is determined by the parameters of agent behaviour and its distribution. In this section we look at the number of steps taken to achieve this spread.

Figure 4,9(a) shows the histogram of the spread of innovation for a particular case. For both the global (upper component) and local (lower component) cascades a histogram of the number of time steps it took to reach the maximum percolations is shown in Figures 4.9(b) and 4.9(c) respectively.

Figure 4.9: (a) Histogram of the final spread of innovation from 1000 runs of the model for the Financial & Professional Services industry (α=0.2, ε=1, S_0=24, β_a=0.25, β_s=1). Histogram of the number of time steps to the maximum percolation of each run of the model which resulted in (b) a local cascade and (c) a global cascade.

(a)

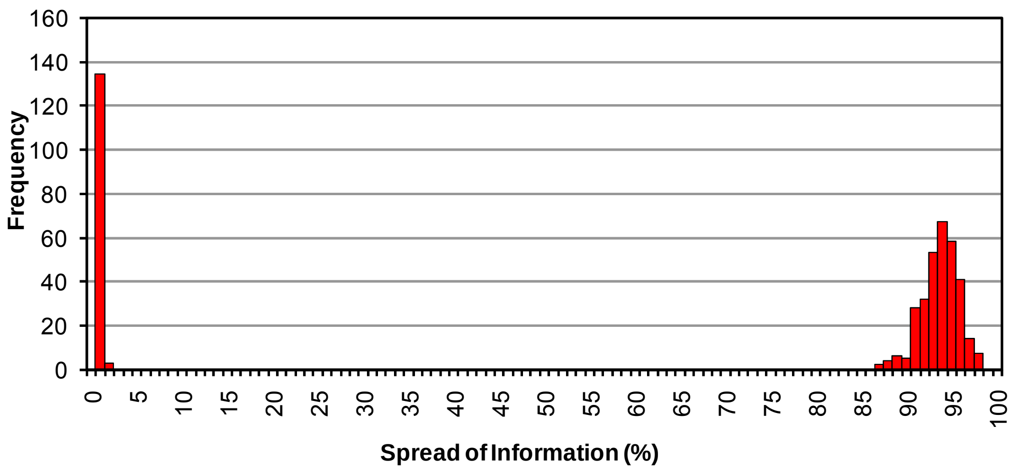

(b) (c)

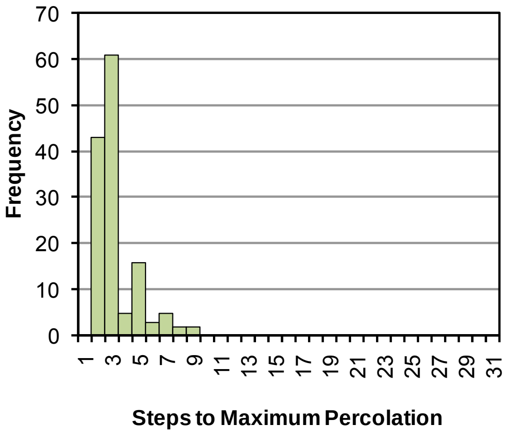

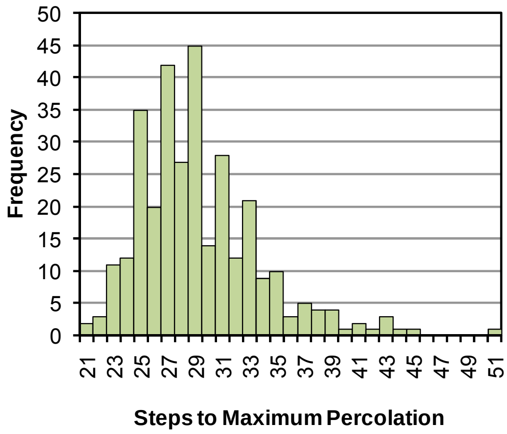

Naturally global cascades require more time steps than the local ones but the initial pattern of both is within the same range. It is only when innovation continues to spread after 4 to 10 steps that we can be sure that there will be a cascade of some size. The model does not directly identify the size of a time step but a natural interpretation would be around three months. This is a reasonable time for an organisation to identify and implement an innovation. So it takes at least a year and maybe more than 2 years before we can be sure that an innovation will not be generally successful. Moreover, it will take between 4 and 7 years for an innovation to reach its maximum extent on this basis.

Figures 4.10 and 4.11 show the upper and lower boundaries for local and global cascades from 1000 runs of the model. They show that there is a wide range of potential behaviour, with spreads of innovation to between 10 and 70 per cent of organisations within 15 time steps in global cascades. For local cascades, it is clear that some innovations peter out after 3 steps, while others take longer.

Figure 4.10: Upper and lower boundaries for the percolations across 310 of the 1000 runs which were global cascades (α=0.4, ε=1, S_0=24, β_a=0.25, β_s=1).

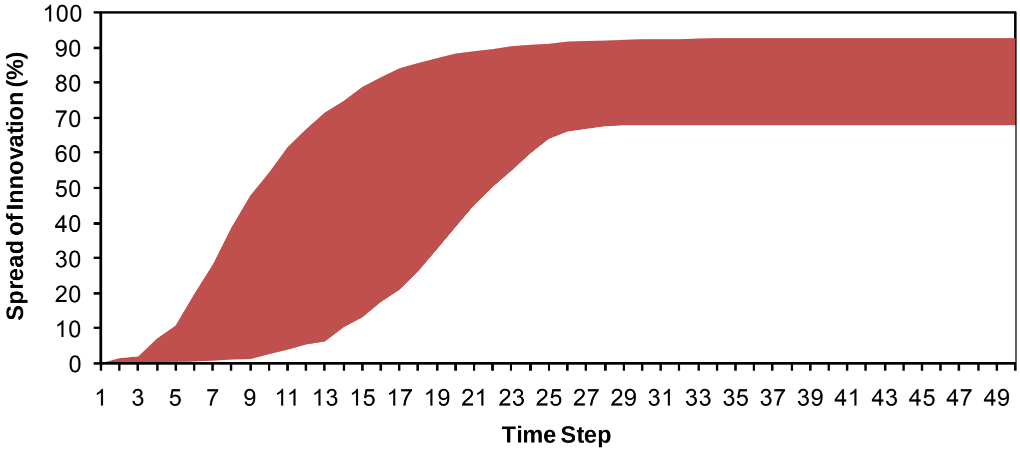

Figure 4.11: Upper and lower boundaries for the percolations across 71 of the 1000 runs which were local cascades (α=0.4, ε=1, S_0=24, β_a=0.25, β_s=1).

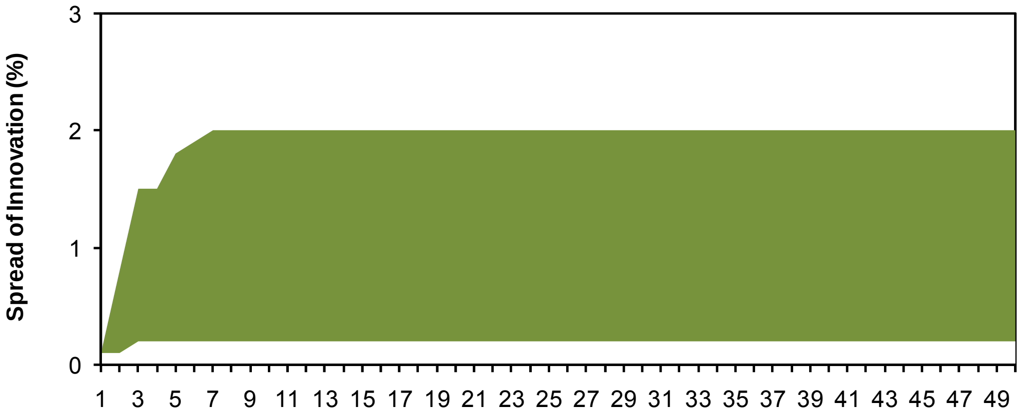

Finally, we look at the relationship between the parameters of the model and the time steps to a cascade. We concentrate on the results for the Financial and Professional Services industry since this demonstrated the most pronounced variations in behaviour.

Figure 4.12(a) shows the average percolation of the industry's global cascade for various parameterisations. As copying becomes less effective, the total spread of innovation falls, and this is also true when the density of the links in the network are reduced (S_0 decreases). It also decreases as the absorptive capacity, β_a, is reduced.

Figure 4.12: Analysis of the average properties of global cascades arising from the 1000 runs of the model for the Financial & Professional Services industry (ε=1, S_0=24, β_s=1) for β_a=0.25 and β_a=0.75 and (ε=1, S_0=16, β_s=1 β_a=0.25). a) Average spread of the innovation, (b) Average time taken to reach maximum percolation and (c) average number of organisations that adopt the innovation per time step.

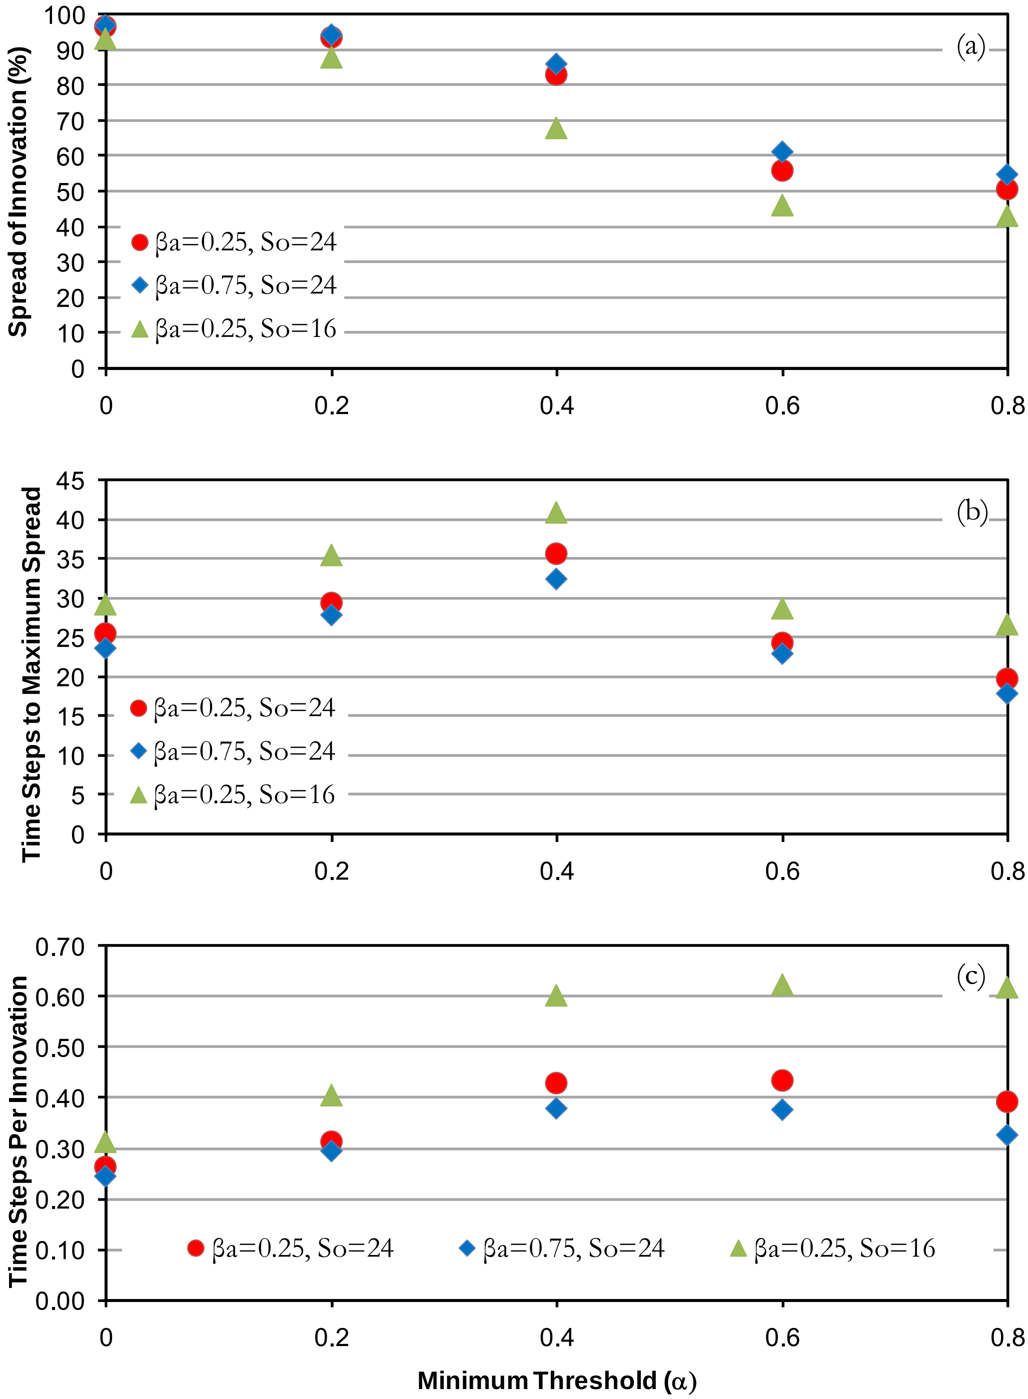

Figure 4.12(b) shows the corresponding average number of time steps taken for the innovation to reach its maximum (similar distributions for the local cascade are also available but the averages are less meaningful because the distributions are asymmetric). It indicates that as copying becomes more difficult or absorptive capacity falls, or the network is less dense, it takes more time to reach the maximum spread of innovation. However, once there are further restrictions on the ability to percolate, both the extent of the spread and the number of steps are both reduced – there is simply less effective innovation.

This transition occurs because, as it is made easier for an innovation to spread, this initially opens up portions of the network that were previously inaccessible. When these bottlenecks are removed the additional organisations that the innovation spreads to are further away from its origin and so the time taken to reach these groups adds a few additional time steps to the model before the model reaches its maximum.

Figure 4.12(c) shows the ratio of (b) and (a), or in other words the time steps per effective innovation. It shows that this rate appears to be more strongly influenced by the density of the links rather than the behavioural rules of the agents themselves.

5 Conclusion and Implications for Policy

This model offers a way of thinking about innovation based on dissemination across a network where the behavioural rules of innovating organisations also have a bearing on the outcome. This seems an entirely plausible way to approach this problem, though even so there are many simplifications around the type of innovation under consideration, let alone the way the innovation comes into being in the first place.

In spite of the considerable simplifications of this model, there still emerge many interacting parameters. Network structure, copying behaviour, and the willingness to exchange interact in different ways according to how they are varied.

Moreover, we are also able to show that the distribution of these parameters across groups of agents also matters. The ability to calibrate these distributions to survey results in a region of the UK enables us to investigate the importance of the distributions as well as the average values of the parameters.

The survey results for the various industries examined do not show very large differences in average behaviour for most of industries – health sciences appears to be an exception. However, when these differences interact across the various parameters, it changes the outcomes very significantly.

In addition, our results also show that the ability to distinguish between a successful, global, cascade of innovation at the outset is very limited. Whether a cascade emerges or not is not predictable at the outset.

These two results, taken together, mean that it is impossible to design policy levers with any certainty as to how they will be effective.. We can see that industry surveys with similar results for innovation behaviour and attitudes can nonetheless produce very variable outcomes. Moreover, initial success cannot predict the final outcome.

Classically, policy has concentrated on identifying innovation that ought to be supported, and on creating sector support groups. In the UK, the latest policy initiative has been the establishment of NESTA (National Endowment for Science, Technology and the Arts). A recent pronouncement suggested setting up new sector groups, to look at different ways innovation takes place in 'non-traditional' sectors. How this is likely to help change the parameters of the kinds of behaviour which generate innovation take-up is not at all clear.

However, the models do suggest that there are ways to help maximise the likelihood of innovation spreading. The model shows that:

- The willingness to exchange innovation is a major driver
- Well-connected but not too dense networks help
- A willingness to copy others helps

Policy can help create networks by events and the support of the widest possible networks. Note that these need to be more broadly focused that on just one sector. Policy to support the willingness to absorb innovation is much more difficult. The policy instinct is to provide grants or tax breaks. However, there is no guarantee these would support the sort of innovation that organisations can readily use, which may be hard to identify in policy terms.

The conclusion can only be that a continual policy of supporting test and try which does not worry too much about an ex ante analysis of the potential for success is the only option.

Appendix I

Calculating the Secrecy Index

Calculating a secrecy index for each firm was done by the following formula:

$$\sum_k A_k \times \left(\sum_i A_k \times \sum i\right) \div \sum_k \left(3 \times \sum_i A_{kk} \times \sum i\right)$$

where i is a firm within a particular industry,

k is an action to protect an innovation,

A_k is a dummy variable that can take values between zero and three. Zero means the firm did not undertake an action while three means the firm attaches high importance to the action,

A_{kk} is a dummy variable that is zero when a firm's value for A_k is between 0 and 2, and takes the value 1 when a firms value for A_k is 3.

Each method of protecting an innovation is weighted by its revealed importance to the sector, where revealed importance is the percentage of firms saying the method was very importing to them. A firm is then given an innovation index, which is the sum of the weights of the methods of protection it did undertake multiplied by the importance the firm put on each method (between 1 and 3). This is then divided by the sum of all the weights multiplied by three. The result is a firm can be given an index between zero and one. A score of one indicates a firm that uses every method of protecting innovation that the survey asked about. A score of zero indicates a firm that used none of the methods of protecting innovation that we asked about.

Appendix II

Calculating the Absorptive Index

Calculating an absorptive index for each firm was done by the following formula:

$$\sum_m A_m \times \left(\sum_i A_m \times \sum i\right) \div \sum_m \left(\sum_i A_m \times \sum i\right)$$

Where i is a firm within a particular industry,

m is an activity conducive to absorbing an innovation,

A_m is a dummy variable equal to one when a firm performed activity m and zero when it did not.

Each method of absorbing innovation is weighted by its revealed importance to a sector. This weight is effectively the percentage of firms in that industry who undertook that method of absorption. A firm is then given an absorptive index which is the sum of the weights of those actions it did undertake, divided by the sum of all the actions weights. A firm can therefore be given an absorptive index between zero and one. Zero indicates a firm that undertook none of the methods of absorbing innovations that the survey identified. One indicates a firm that engaged in every method of absorbing innovation that the survey asked about.

Appendix III

Calculating the Innovation Index

Calculating an innovation index for each firm was done by the following formula:

$$\sum_j A_j \times \left(\sum_i A_j \times \sum i\right) \div \sum_j \left(\sum_i A_j \times \sum i\right)$$

Where i is a firm within a particular industry,

j is a type of activity conducive to innovation,

A_j is a dummy variable equal to one when a firm performed activity j and zero when it did not.

Each innovation is weighted by its revealed importance to a sector. This weight is effectively the percentage of firms in that industry who undertook that innovation. A firm is then given an innovation index which is the sum of the weights of those innovations it did undertake, divided by the sum of all the innovation weights. A firm can therefore be given an innovation index between 0 and 1 with a lower value indicating a lower propensity to innovate.