

Extremism Propagation in Social Networks with Hubs

Daniel W. Franks^{*†}, Jason Noble^{**}, Peter Kaufmann^{***}, and Sigrid Stagl^{***}

df525@york.ac.uk, dr.Jason.Noble@googlemail.com, [P.Kaufmann | S.Stagl]@sussex.ac.uk

* York Centre for Complex Systems Analysis (YCCSA), Department of Biology &

Department of Computer Science. The University of York, York, YO10 5YW, UK

** School of Electronics and Computer Science, University of Southampton, Southampton,

SO17 1BJ, UK

*** SPRU - Science and Technology Policy Research, University of Sussex, Brighton, BN1

9QE, UK

† Author to whom correspondence should be sent. Tel: +44 (0)1904 328648

Abstract

One aspect of opinion change that has been of academic interest is the impact of people with extreme opinions (extremists) on opinion dynamics. Amblard and Deffuant (2004) used an agent-based model to study the role of small-world social network topologies on general opinion change in the presence of extremists. They found that opinion convergence to a single extreme occurs only when the average number of network connections for each individual is extremely high. Here, we extend the model to examine the effect of positively skewed degree-distributions, in addition to small-world structures, on the types of opinion convergence that occur in the presence of extremists. We also examine what happens when extremist opinions are located on the well-connected nodes (hubs) created by the positively skewed distribution. We find that a positively skewed network topology encourages opinion convergence on a single extreme under a wider range of conditions than topologies whose degree distributions were not skewed. The importance of social position for social influence is highlighted by the result that, when positive extremists are placed on hubs, all population convergence is to the positive extreme even when there are twice as many negative extremists. Thus, our results have shown the importance of considering a positively skewed degree distribution, and in particular network hubs and social position, when examining extremist transmission.

Keywords: Social Networks, Scale-Free, Small World, Extremism, Opinion Change

1. Introduction

People hold opinions about an almost countless number of topics. From religion, politics, literature and moral values, to entertainment, hair products, and the lives of celebrities, most people will have an opinion on the matter. These opinions can be shaped by individual reflection, but are in many cases greatly influenced by social context: if your friends all prefer Dan Brown to Dostoevsky, then you are likely to do so as well. Dual inheritance theory (Boyd & Richerson, 1985) tells us that human behaviour cannot be understood without considering both our genetic and cultural inheritances. But our cultural inheritance is not absorbed at random or from all possible sources: opinion formation is embedded in social interaction (Wood, 2000). In other words, what you think depends on who you talk to, and whose opinions you respect. It follows that understanding the process of social influence and the social structures that govern human interaction will be central to an understanding of the dynamics of opinion change. Opinion change across a whole population, of course, constitutes a change in cultural norms, and thus the dual inheritance program will remain incomplete without an improved understanding of human social structures and their effects. Given that opinions influence behaviour, a full account of human behaviour, both adaptive and maladaptive, likewise depends on a theory of how opinion change is rooted in social interaction.

One aspect of opinion change that has been of academic interest is the impact of extremists on opinion dynamics. Extremists are people with marginal opinions that are far away from ambivalence. The multidisciplinary innovation diffusion literature has produced numerous case studies where an initially small minority of extremists spread their opinion to a majority

of the population. Examples include drug use, new cropping techniques, and the introduction of family planning practices (Rogers 1995). The spread of extreme opinion in Germany in the 1930s, and in Rwanda in the 1990s, stand out as salient and horrific examples. On the other hand, there are extremists whose opinions (thankfully) never became generally accepted, such as those of the KuKlux Klan. All minority rights activists, such as those campaigning for womens' right to vote or those involved in the American civil rights movement, were considered at some point to be extremists, although their views became mainstream.

Societal structure is important to opinion change, and can be conveniently represented as a social network. In a social network, nodes represent individuals and connections represent interactions between those individuals. Many studies have indeed shown that different types of social interactions, such as scientific collaborations and sexual contacts, can be understood in terms of networks (Newman, 2003). Evidence has shown that social networks have distinctive properties (Newman & Park, 2003). Most prominently they are usually characterized as small-world and scale-free networks. In a small-world network, each node can be reached by relatively few steps, due to the existence of long-distance connections bridging different areas of the network. Scale-free networks have a degree distribution that is positively skewed to the right in a manner similar to a power-law distribution. This results in some nodes possessing a far greater number of connections than others. However, scale-free is a theoretical term that needs to be used with care (especially when dealing with networks of just hundreds of thousands of nodes). Here we refer to our analogous degree distribution in our relatively small networks as being positively skewed. It may be useful to carefully draw some appropriate analogies with the properties of scale-free networks. Our own data (taken from an unrelated study and shown as an example) collected on farmer social networks in six countries, clearly shows a positively skewed degree distribution (Figure 1). Thus, when

studying the effects of social network topologies it may be important to capture their small-world and positively skewed degree distribution structures.

There have been numerous attempts at building models of opinion change. Many models consider binary opinions (e.g., Latané and Nowak, 1997; Kacpersky and Holyst, 2000). Because the opinions in these models are only binary representations they do not distinguish between hard-core extremists and varying levels of tacit supporters and non-supporters. Axelrod (1986) and Schelling (1978) evaluated opinion change using simple lattices. They both showed that opinions tend to polarize (i.e., multiple opinions coexist). However, the opinions were modelled as discrete, and thus extremists were not present. As such, these can only be used to examine the spread of an opinion to become the minority or majority opinion. The bounded confidence model (Krause, 2000; Dittmer, 2001) relaxes the assumption of binary opinions. Instead, opinions are represented as points on an opinion continuum (e.g., anywhere between -1 and 1). Interactions become non-linear in that agents influence each other only if the distance between their opinions is below a threshold. The relative agreement model is an extension of the bounded confidence model. In the relative agreement model the level of influence between agents is governed by the distance between opinions and the certainty of the influencing agent, rather than some predefined level of influence. The model captures the following ideas: (a) agents with radically different opinions are unlikely to influence each other; (b) uncertain agents are more susceptible to social influence than agents that are certain; and; (c) agents that are certain are more influential than uncertain agents. The assumptions of the relative agreement model have been shown to be supported well by laboratory experiments (Deffuant et al., 2002), such as those dealing with opinion radicalization (for one of the first examples see, Moscovici and Lécuyer, 1972). Whereas these models allow for influence based on persuasion (attractive forces), they do not allow for

conflicts of opinion to drive opinions further apart (repulsive forces). The meta-contrast model (Salzarulo, 2006) allows for both attractive and repulsive influence. Thus, the model allows for the possibility that the final opinions can become more extreme than the most extreme initial opinion.

Amblard and Deffuant (2004) studied the role of the social network topology on extremist transmission (see also, Deffuant, 2006). They developed an agent-based model of social influence using the relative agreement model, and examined the type of opinion convergence that occurred under various conditions with various small-world topologies (as opposed to a fully connected network). Extremists were represented as agents with opinions at the extreme ends of the opinion spectrum (see below for a more technical description). They found that, for the small-world topology, a critical level of average network degree was needed for the population to converge on a single extreme opinion. However, the average network degree needed to be extremely high for this to occur; typically requiring an average degree of around 64 or even higher for lower levels of random connections.

In this paper we further extend the model (Amblard and Deffuant, 2004) to examine the effect of a positively-skewed-degree-distribution structure on the types of opinion convergence that occur in the presence of extremists¹. We also examine what happens when extremist opinions are located on well-connected nodes.

¹ Weisbuch et al. (2005) briefly look at scale-free structures for the relative agreement model, but do not examine their effect in any depth, and also do not explore social position or hubs.

2. The Model

Our agent-based model consisted of one thousand agents embedded in a social network. Each agent i was characterized by two attributes: an opinion x_i and level of uncertainty u_i . Social influence was assumed to take place between all connected individuals in a random order each time-step. In all cases simulations were run for 2000 time-steps, after which all simulations found a stable-state.

2.1. Initialization

The population included N individuals of which there were $p_e.N$ extremists and $(1-p_e).N$ non-extremists. p_+ and p_- were the proportions of positive and negative extremists ($p_e = p_+ + p_-$). Thus, the population was initialized with $p_+.N$ positive extremists and $p_-.N$ negative extremists. When we talk of extremists as being positive and negative we are simply referring to their position on an abstract opinion continuum (-1 or +1) without any connotations as to the social value of their opinions (i.e., a negative extremist need not hold a "negative" opinion in the colloquial sense). Non-extremist opinions were initially drawn from a uniform distribution between -1 and 1, and the initial uncertainty of non-extremists was set to U . Extremists were assumed to be those agents with opinions located at the extremes of the opinion distribution x_i (i.e., -1 or 1). They were also assumed to be more confident (i.e., have a lower value of u_i) than non-extremist individuals with an initial uncertainty of u_e (where $u_e < U$). For all simulation runs $p_e = 0.1$ (Deffuant, 2006) and we assumed that equal proportion of positive and negative extremists in the population ($p_+ = p_- = p_e/2$).

2.2. The Relative Agreement Model

The relative agreement model captures social influence between two agents (Deffuant et al., 2002). In this model, agents' opinions can be visualized as segments with boundaries defined as $x_i - u_i$ and $x_i + u_i$. In other words, the centre of the segment is defined by x_i and the level of uncertainty defines the distance of the opinion edges either side of x_i . Agent i can influence agent j , bringing agent j 's opinion closer to agent i 's opinion, if their opinion segments overlap (i.e., they have at least *some* common ground). The intensity of the influence depends on the amount of overlap between the opinion segments; the higher the overlap the higher the influence. The model also captures the notion that agents that are more certain about their opinion are more influential. As an agent is influenced, its level of uncertainty is slightly reduced. A formal statement of a relative agreement between two agents, i and j , follows.

The width of the overlap between opinion segments is given by:

$$h_{ij} = \min(x_i + u_i, x_j + u_j) - \max(x_i - u_i, x_j - u_j) \quad (1)$$

Thus, the non-overlapping width is:

$$2u_i - h_{ij} \quad (2)$$

The agreement is defined by the overlap minus the non-overlap:

$$h_{ij} - (2u_i - h_{ij}) = 2(h_{ij} - u_i) \quad (3)$$

The relative agreement is the agreement divided by the length of agent i 's segment:

$$\frac{2(h_{ij} - u_i)}{2u_i} = \frac{h_{ij}}{u_i} - 1 \quad (4)$$

If $h_{ij} \leq u_i$, then there is no influence. This corresponds to the idea that agent j 's opinion segment must intersect agent i 's actual opinion (at the centre of its segment) for agent i to consider agent j 's opinion. Otherwise, the modifications of agent j 's opinion x_j and uncertainty u_j as a result of the interaction with agent i multiplied by the relative agreement:

$$x_j = x_j + \mu \left(\frac{h_{ij}}{u_i} - 1 \right) (x_i - x_j), \quad (5)$$

$$u_j = u_j + \mu \left(\frac{h_{ij}}{u_i} - 1 \right) (u_i - u_j), \quad (6)$$

where μ governs the influence intensity for all interactions. A high value of μ results in quick opinion convergence between neighbouring agents.

2.3. Constructing the Social Network

Networks were constructed in a manner that allowed us to vary the extent to which they possessed small-world and positively-skewed-degree-distribution characteristics (Noble et al., 2004). The average number of connections for each agent is denoted k . Each agent's social position was represented as a network node, arranged along a one-dimensional ring lattice. The distance between nodes on this lattice defines the agent's local neighbourhood; for a neighbourhood of size n ($n = 2k$) this would consist of the nodes contained within a distance

of k either side of the focal node. There were two axes on which the layouts of our networks could vary. The first axis of variation allowed us to control the frequency of random long-range connections. Thus, we used a parameter R that varied the probability with which a connection would be either local (i.e., from the neighbourhood) or random. A local connection meant a connection to a randomly chosen node within the local neighbourhood. A random connection meant a connection to a randomly chosen node from anywhere else on the network. When $R = 1$, all links were random, when $R = 0$ all links were local. Small intermediate values of R create a network with small-world properties (Watts & Strogatz, 1998).

The second axis of variation allowed us to control the level to which the network degree distribution was positively skewed (Figure 1). This was done using the parameter P , which allowed us to vary the level of preferential attachment to well-connected nodes. When $P = 0$ each node was equally likely to be selected for a connection. As the value of P increased connections were made in an increasingly preferential way, and well-connected nodes were more likely to receive a connection than poorly-connected nodes. The value $(d + \sigma)^P$ was calculated for each eligible node, where d is the node's degree and σ is a small positive value (0.1 in this model) that ensures all nodes have a chance of selection. The node selected to receive a connection was chosen by roulette-wheel selection over these values. Thus, the parameter P is a preferential exponent, governing the strength of the bias towards connecting to well-connected nodes (see also, Barabási & Albert, 1999).

[FIGURE 1 HERE]

To construct the network, nodes were sorted in a random order and given one initial connection, selected according to the current values of R and P . This procedure helped to ensure that all nodes were part of the network. Then connections were allocated until K average connections were reached. The values of R and P were used in the selection of both source and destination nodes for each connection.

2.4. Convergence Indicator

The model typically results in three types of opinion convergence: ‘central convergence’, ‘both-extremes convergence’, and ‘single convergence’. Central convergence occurs when extremists have little influence on the general opinion. Both-extremes convergence occurs when extremists with opposing opinions both influence general opinion enough to cause opinions to cluster around both extremes. Single extreme convergence occurs when general opinion converges upon a single extreme opinion. Deffuant et al. (2002) found that the convergence type exhibited by a simulation run can be conveniently found using an indicator y . After a population has converged, y is calculated as follows:

$$y = p'_+{}^2 + p'_-{}^2 \quad (7)$$

where p'_+ and p'_- are the proportions of initial non-extremists that became extremists to the positive and negative extreme respectively.

y indicates the type of convergence as follows:

- No extreme convergence: $y = 0$

- Double extreme convergence: $y = 0.5$
- Single extreme convergence: $y = 1.0$

Intermediate values of y correspond to intermediates of the above situations.

3. Results

Amblard and Deffuant (2004) explored opinion convergence on their small-world network for variations of k (the average degree) and R (the proportion of non-local connections). We follow this approach for our study. When varying k and R the following parameters are kept constant: $U = 1.8$, $p_e = 0.1$, $\mu = 0.1$. Thus, we are studying opinion formation where non-extremists begin with arbitrary and uncertain opinions. Lower values of U could stop the population all converging on extreme opinions, and using an unequal ratio of positive and negative extremists would aid single extreme convergence. However, these parameters have been exhaustively studied elsewhere (e.g., Deffuant et al., 2002; Amblard and Deffuant, 2004) and do not contribute to the argument made in this paper. First, we replicate the work by Amblard and Deffuant by studying a small-world network. Second, we extend the model to allow for positively skewed degree distributions to be constructed. Finally, we examine the effect of placing extremists on well-connected network nodes. In the forthcoming graphs (Figure 2 to Figure 7) the waves and occasional spikes of the lines are simply noise, and the graphs would be smoother if we averaged the values over thousands of trials. We were limited to 50 replications due to the computational time required for the runs.

[FIGURE 2 HERE]

3.1. *Small-World*

Small-world networks were constructed by varying R between 0 and 1 ($P = 0$). Figure 2 shows the type of convergence found for variations of k and R . Notice that for single extreme convergence to occur regularly (e.g., $y > 0.8$), the network must have a relatively high number of random connections, and the average degree must be extremely high (typically $k > 60$). Single extreme convergence is always stable as the network tends to saturate in this state. This result confirms those of Amblard and Deffuant (2004).

[FIGURES 3 & 4 HERE]

3.2. *Positively Skewed Degree Distribution and Small-World*

Small-world networks were constructed by varying R between 0.2 and 1. When $P = 1$, we found that it gives positively skewed degree distributions that qualitatively match our farmer social networks data. Note that our intention here is to conduct a theoretical study, it is not to accurately model the farmer data. We simply show both distributions to illustrate that both network degree distributions are positively skewed, with many individuals connected to few individuals, and few individuals connected to many individuals. Figure 3 shows the type of convergence found for variations of k and R . For single extreme convergence to generally occur, the network does not need an average degree as high as in the case where the network degree is not skewed, and the proportion of random connections does not need to be as high. Influence by extremists is higher for almost all data-points in the case with the skewed

distribution. Figure 4 shows the type of convergence found for variations of k and P . Convergence to a single extreme occurs readily for high values of k and P .

[FIGURES 5, 6 & 7 HERE]

3.3. *Extremists at Network Hubs*

Networks with positively skewed degree distributions, like scale-free networks, produce a relatively limited number of well-connected nodes, or hubs (and random network also produce hubs, but a limited number of poorly connected nodes). We introduced a new parameter λ to control the proportion of positive extremists placed at these hubs. Thus, once the network had been constructed, $p_+ \cdot N \cdot \lambda$ positive extremists exchanged their network position with the agents on the best-connected nodes. Figure 5 shows the type of convergence found for variations of k and R . The results are similar to the case with the skewed degree distribution without positive extremists on well-connected nodes, although for values of $R < 0.2$ a lower average degree is needed for single-extreme convergence.

To examine the effect of placing positive extremists on key nodes on the trajectory of opinion change, we plotted the proportion of non-extremists that became positive extremists for networks with positively skewed degree distributions, with and without positive extremists at hubs (Figure 6). We found that placing positive extremists at network hubs greatly increased the proportion of non-extremists that become positive extremists. Under some conditions (roughly $k > 22$) most of the population (more than 80%) became positive extremists.

Figure 7 shows that even when negative extremists are twice as prevalent in the population, opinion convergence to the positive extreme can regularly occur if positive extremists are placed on network hubs. We reproduced the conditions for Figure 7 without placing positive extremists at key hubs. The results showed that, under all conditions, the population never converged on the positive extreme (no graph shown).

4. Discussion

A positively skewed degree distribution increases opinion convergence towards extremes, and encourages opinion convergence on a single extreme under a wider range of conditions than topologies that were not skewed in their degree distribution. For single extreme convergence to occur for the small-world model, Amblard and Deffuant (2004) required an extremely large number of average connections. Such a large average number of connections is unlikely to exist in a real-world social network. Constructing a positively skewed degree distribution reduces the number of connections needed for single-extreme convergence to a more realistic number of average social network connections. Another effect of the skewed distribution is that there is little difference between a small-world network with few random connections (e.g., $R = 0.25$) and a random network ($R = 1$).

Why does a positively skewed degree distribution increase extremism transmission in a social network? We suspect that the creation of some well-connected nodes (hubs), along with the creation of many poorly connected individuals, is responsible. Extremists located at these key nodes are connected to more agents, and more parts on the network. Thus, their influence is more frequent and wider-reaching, making it less likely that certain parts of the network remain socially isolated from the extremist's opinion. In the default case, most of the time

extremists will not be located at hubs. However, extremists take advantage of non-extremists at hubs, as a means of transmitting their opinion. Non-extremists are initially less confident about their opinion than extremists. Thus, extremists have more influence over non-extremists than non-extremists have over extremists (and more influence than non-extremists have over each other). Once an extremist has spent time influencing a non-extremist at a hub, the former non-extremist will transmit the opinion throughout the network. Because agents at hubs are well-connected, it is likely that they are connected to an extremist. In the real world, extremists might purposely attempt to influence people in high status positions.

Valente (1995) investigated some effects of propagating innovations first to more ‘strategic’ nodes of the network, showing an individual’s network position can affect innovation diffusion in a social network. Here, we examined the effect of extreme opinions at high social positions. By placing positive extremists at the well-connected nodes (i.e., nodes to the tail-end of the degree distribution) we found that when the network converges on a single extreme, it is always to the positive extreme. This is because when positive extremists are placed on key nodes they are able to directly influence more agents (as they would have more direct connections) and have a higher chance of having a random connection that allows them to influence a different social network neighbourhood. The importance of social position for social influence was highlighted by the result that, when positive extremists are placed on hubs, all population convergence is to the positive extreme even when there are twice as many negative extremists. Thus, the position of an opinion in the social network is more important than the initial proportion of individuals with the opinion. The network does not have to have a power-law distribution for this result to stand. However, the network must have multiple hubs and social networks are known to produce such hubs. In the real world, certain extremists are likely to fight for high-status network positions in order to preach their

opinions. For example, people with extreme political opinions might actively attempt to seek higher social status or widen their social network with the intention of spreading their opinion.

These findings are of relevance for applications in social science, biology, and computer science. In social science, one can think of multiple applications where the structure of networks is of relevance. We have primarily shown how certain social structures can contribute to changing a population's values, attitudes and behaviour. Religious extremism is one of many examples where an extreme opinion exists, and thus has the possibility of percolating throughout a sub-community such as the one modelled here. We have shown that such opinions are less likely to pervade the social network if well connected, typically high status, people are resistant to the opinion. The findings are also useful for studies of innovation diffusion, and market activities where geographic actor interactions are explicitly taken into account. In biology, the social network structure can have a large affect on social learning in animals. For example, in a social learning model that did not explicitly consider network structure, Noble and Franks (2002) found that who learns from whom affected the efficiency of different social learning mechanisms. Given our findings, that agents' social positions greatly affect their influence, these results might be emphasized. In computer science, knowing more about social networks should allow us to construct better software for monitoring, shaping, and exploiting existing human networks. There could also be feedback into the design of large-scale artificial multi-agent systems, in which the agents necessarily communicate over a social network.

Our results have shown the importance of considering a positively skewed degree distribution when examining extremist transmission. Future work will look at the effect of modelling

loosely coupled sub-communities, network formation, network growth, and coevolving network structure with opinions. Such studies will help us to further understand opinion change and extremism transmission. An interesting future question that this model could be used to explore is: if most social networks are small-world and have positively skewed degree distributions, then why do we see the persistence of extreme opinion polarization in populations? It is clear that social behaviours (and thus human behaviour in general) cannot be fully understood without considering their underlying network structures. Agent-based models can help us look for underlying principles that allow us to understand social behaviours in terms of social networks.

Acknowledgements

We thank two anonymous referees for their comments and suggestions. The model presented in this paper was developed while D.W.F. was at SPRU (The University of Sussex) working in the IDARI project, financed under the FP5 Quality of Life and Management of Living Resources, Key Action 5.

Figure Captions

Figure 1: The qualitative similarity between degree distributions for (left) a real-world farmer social network in Estonia ($N = 170$, $K = 7.07$); and (right) a simulated social network in our model, with highly preferential attachment ($N = 1000$, $k = 7.07$, $R = 0.2$, $P = 1.0$). Note the similarities in the general shape of each graph. They both show a skewed network degree distribution where many nodes have few connections and few nodes have many connections.

Figure 2: (left) Contour plot of the mean opinion convergence indicator values, y , for variants of k and R , averaged over 50 runs. (right) The standard deviation of y . The graph shows the result of varying the network topology between only local connections, and only random connections; with small-world topologies existing between the extremes of R . The standard deviation for some mid-range values of y indicates that some trials may have resulted in single extreme convergence, and some may have reached central convergence.

Figure 3: The mean opinion convergence indicator values, y , averaged over 50 runs for variants of k and R , when the degree distribution is positively skewed ($P = 1$). The graph shows the result of varying the network topology between only local connections, and only random connections; with small-world topologies existing between the extremes of R . Notice that for single extreme convergence to generally occur, the network does not need an average degree as high as in the case where the network degree is not skewed. Influence by extremists is higher for almost all data-points in the case with the positively skewed degree distribution.

Figure 4: The mean opinion convergence indicator values, y , averaged over 50 runs for variants of k and P , when $R = 0.5$. The graph (left) shows the result of varying the level of preferential attachment to well-connected nodes; with positively skewed degree distributions existing for higher values of P . Note that convergence to a single extreme occurs readily for high values of k and P . The standard deviation is shown on the right.

Figure 5: The mean opinion convergence indicator values, y , averaged over 50 runs, for variants of k and R , when the degree distribution is positively skewed ($P = 1$) and positive extremists are placed on the most well-connected nodes ($\lambda = 1$). The graph shows the result of varying the network topology between only local

connections, and only random connections; with small-world topologies existing between the extremes of R . Notice that the graph is a smoothed-out version of the graph in figure 3. Influence by extremists is higher for almost all data-points in the case with the positively skewed degree distribution. The standard deviation is close to zero for all data-points.

Figure 6: The mean proportion of initial non-extremists that become an extremist to the positive extent, averaged over 50 runs. (left) When the degree distribution is positively skewed ($P = 1, \lambda = 0$), (right) when the degree distribution is positively skewed, and positive extremists are placed on the most well-connected nodes ($P = 1, \lambda = 1$). The graph shows the result of varying the network topology between only local connections, and only random connections; with small-world topologies existing between the extremes of R .

Figure 7: The mean proportion of initial non-extremists that become an extremist to the positive extent, averaged over 50 runs, when the degree distribution is positively skewed ($P = 1$) and positive extremists are placed on the most well-connected nodes ($\lambda = 1$). In this case, there were twice as many negative extremists as positive extremists. The graph shows the result of varying the network topology between only local connections, and only random connections; with small-world topologies existing between the extremes of R . Note that despite there being twice as many negative extremists; the population converges to the positive extreme under a range of conditions.

References

- Amblard, F., & Deffuant, G. (2004). The role of network topology on extremism propagation with the relative agreement opinion dynamics. *Physica A*, 343, pp. 725-738
- Axelrod, R. (1986). An evolutionary approach to norms. *American Political Science Review*, 80, pp. 1095-1111.
- Barabási, A.-L., & Albert, R. (1999). Emergence of scaling in random networks. *Science*, 286, pp. 509–512.
- Boyd, R. & Richerson, P. J. (1985). *Culture and the Evolutionary Process*. Chicago University Press.
- Deffuant, G (2006). Comparing Extremism Propagation Patterns in Continuous Opinion Models, *Journal of Artificial Societies and Social Simulation*. *Journal of Artificial Societies and Social Simulations*, 9(3)
- Deffuant, G., Amblard, F., Weisbuch, G. & Faure, T. (2002). How can extremism prevail? A study based on the relative agreement interaction model. *Journal of Artificial Societies and Social Simulations*, 5(4).
- Dittmer, J.C. (2001). Consensus formation under bounded confidence. *Nonlinear Analysis*, 47, pp. 4615-4621.

- Kacpersky, K. & Holyst, J. (2000). Phase transition as a persistent feature of groups with leaders in models of opinion formation. *Physica A*, 287, pp. 631-643.
- Krause, U. (2000). A discrete nonlinear and non-autonomous model of consensus formation, in S. Elaydi, G. Ladas, J. Popenda and J. Rakowski (Eds.), *Communications in Difference Equations*, Gordon and Breach Science Publishers, pp. 227-236.
- Latane, B. & Nowak, A. (1997). Self-Organizing Social Systems: Necessary and Sufficient Conditions for the Emergence of Clustering, Consolidation, and Continuing Diversity. In: *G.A. Barnett and F.J. Boster (eds.), Progress in Communication Sciences*. Ablex Publishing Corporation, pp. 1-24.
- Moscovici, S. & Lécuyer, R. (1972). Studies on polarization of judgments. *European Journal of Social Psychology*, 2, pp. 221-244
- Newman, M.E.J. & Park, J. (2003). Why social networks are different from other types of networks. *Physical Review E*, 68(3)
- Newman, M.E.J. (2003). The structure and function of complex networks. *SIAM Review*, 45, pp. 167–256.
- Noble, J. and Franks, D. W. (2004) Social learning in a multi-agent system. *Computing and Informatics*, 22(6) pp. 561-574.

Noble, J., Davy, S. & Franks, D.W. (2004). Effects of the topology of social networks on information transmission. *From Animals to Animats 8: Proceedings of the Eighth International Conference on Simulation of Adaptive Behavior*, pp. 395-404.

Rogers, E.M. (1995). *Diffusion of Innovations*. 4th edition, New York: The Free Press

Salzarulo, L. (2006). A Continuous Opinion Dynamics Model Based on the Principle of Meta-Contrast, *Journal of Artificial Societies and Social Simulation*, 9(1)

Schelling, T. (1978). *Micromotives and macro-behavior*. Norton, New York.

Watts, D.J. & Strogatz, S.H. (1998). Collective dynamics of ‘small-world’ networks. *Nature*, 393, pp. 440–442.

Weisbuch, G., Deffuant, G. & Amblard, F. (2005). Persuasion dynamics, *Physica A*, 353, pp.555-575.

Wood, W. (2000). Attitude change: persuasion and social influence, *Annual Review of Psychology*, 51: 539-570.

Valente, T. (1995). *Network Models of the Diffusion of Innovations*, Hampton Press


Figure 1


Figure 2


Figure 3

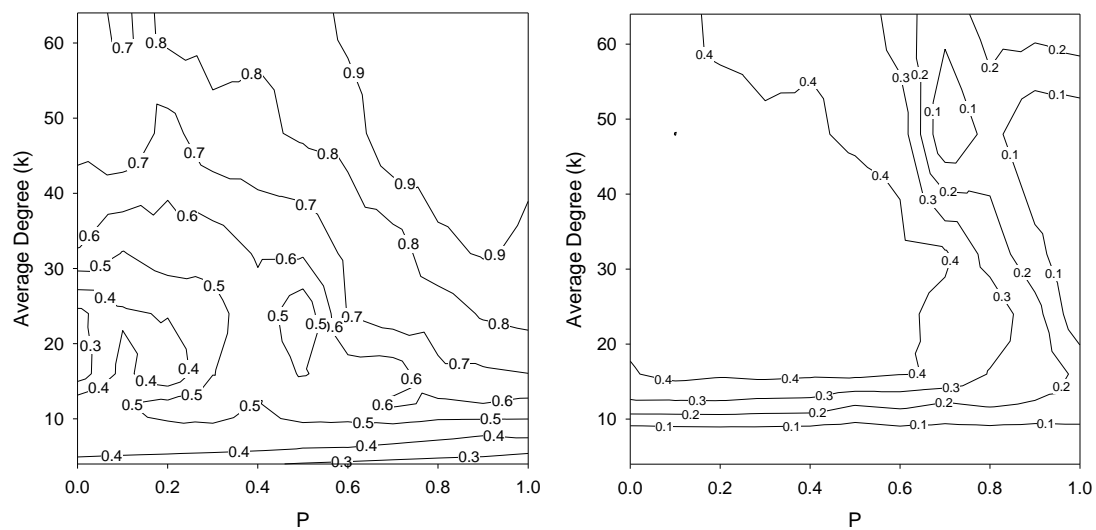
Figure 4


Figure 5

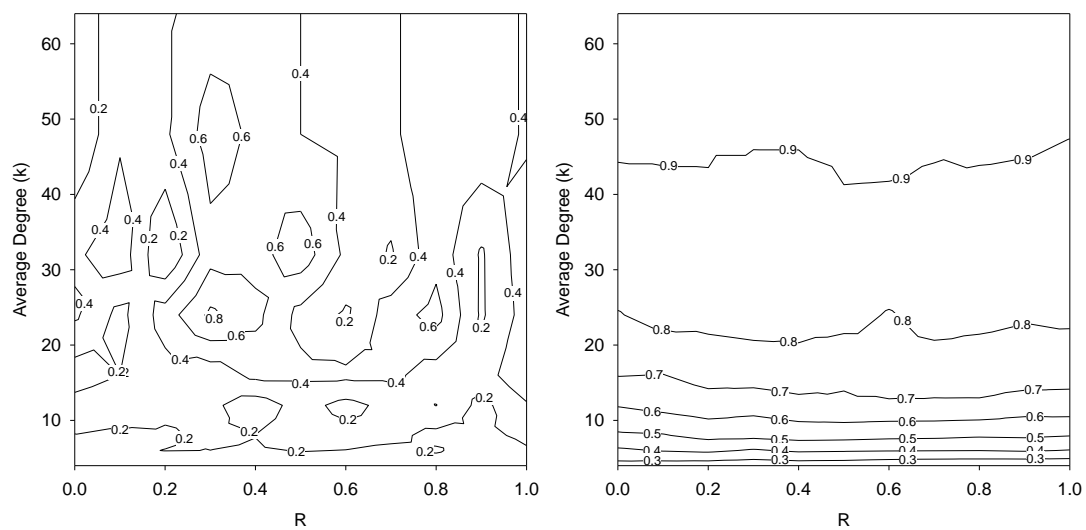

Figure 6

Figure 7

Daniel Franks Bio

Dan Franks is a research fellow in the York Centre for Complex Systems Analysis, in both the Department of Biology and the Department of Computer Science. His research centres on building computational models of ecology and evolution. In particular he is interested in animal social networks (and how best to sample them), and predator-prey interactions (with a focus on warning signals, mimicry, and anti-predatory defenses). Address: York Centre for Complex Systems Analysis, The University of York, York, YO10 5YW, UK. Email: df525@york.ac.uk

Jason Noble Bio

Jason Noble is a research fellow in the Science and Engineering of Natural Systems group, part of the School of Electronics and Computer Science at the University of Southampton. His main research interest is in using individual-based simulations to look at the evolution of social behaviour, particularly communication and social learning. Other interests include network theory, game theory, philosophy of science, and philosophy of mind. Address: School of ECS, University of Southampton, SO17 1BJ, UK. Email: jn2@ecs.soton.ac.uk.