Investigating the evolution of trust across a network utilising agent-based modelling

By: Matthew Oldham & Alessandra Romani

Introduction

We develop an agent-based model that investigates how information flows, and trust evolves, across an evolving network. Specifically, the model starts with a lattice network with 200 agents, who can choose to send high or low quality information to their neighbours. After a given period, agents assess the quality of the information they have received and adjust their trust in their neighbours accordingly. Agents, at given intervals, then decide to maintain the relationship with their existing neighbours or find new neighbours, thus evolving the network. By making agents' trust in their neighbours conditional on a series of parameters, the model shows how the quality of information shapes network structures. We present agents' characteristics and decision rules, discuss the results from the simulation and propose ways in which the model could be extended and applied to social science.

Model Setup

There are five parameters characterizing this model:

- Agents' trust in each of their neighbours;
- Agents' propensity to trust their neighbours;
- Agents' propensity to forgive their neighbours;
- The time lag with which agents verify the quality of information;
- Agents' tolerance of their trust in their neighbours.

All parameters range from 0 to 1, but agents' propensity to trust and to forgive is randomly distributed across the population. This provides a level of heterogeneity amongst the agents. An agents' trust in their neighbours is initiated at 1 at the beginning of the simulation but evolves through time. This latter choice can be justified because agents' level of trust helps determine whether high or low quality information circulates throughout the network—if we were to randomly assign this parameter, agents might end up with values of trust that are so low as to prevent high quality information from flowing effectively. The agents' tolerance is a global parameter, with its relevance discussed in the following section.

Model Dynamics

At the beginning of the simulation, agents gather external information, with a random probability that it is of high quality. Then, agents weigh this information and decide how much of it they are willing to send to their network. The weight is based on the level of trust they have in a given neighbour. The rationale being that, if an agent trusts the given neighbour, they will be more willing to provide their full information.

Once the information is transmitted, and with some lag, agents assess its quality and update their trust accordingly. Specifically, agents compare the information they have received from each of their neighbours to the information that was generated in the entire population—if the quality of the individual information is better than the average information, then trust increases; conversely, if the quality of the individual information is worse than the average information, trust decreases. In assessing the global information, agents potentially discount the information, providing a level of lenience. As for the extent that agents reward or penalize their neighbours, it is dependent on their propensity to trust and to forgive, respectively.

Whenever the simulation gets to the rewiring step, the agents' performance in generating information is compared to the global information, with the top 20 agents being selected as potential new neighbours. Then, agents go through each of their neighbours and assess their trust against the global tolerance parameter, which reflects agents' tolerance for trust—the rationale being that trust levels are a function of an agent's ability to provide high quality information. The tolerance threshold is used to determine whether a link between agents will be maintained: if agents' trust in their neighbours is lower than the threshold, the link between the agent and that neighbour is cut, and vice versa. Knowing how many links they have interrupted, agents go back to the list of top 20 performers and randomly select new links by the same amount. As it was the case at the beginning of the simulation, agents have full trust in their new neighbours and also reset the trust levels with their retained neighbours. The model then continues the interactions per the decision rules stated above, until the next rewiring step.

Results

The model described above provides a great deal of flexibility for testing the dynamics of trust diffusion across a network. To perform an initial assessment, a parameter setting utilizing the settings in Table 1 was performed. The model was also able to capture a great deal of data relating to

both the individual agents and the entire network. The initial analysis, as illustrated in Figure 1 through 4, relates to the trust within the population, the number of links in the network (a proxy for the number of neighbors), the number of switches made by the agents, and the clustering coefficients.

Table 1: Parameter sweep settings

Variable	Settings
Rewire timing	51,101,167, and 250
Tolerance	0, 33%, 66% and 100%
Discount Factor/Benchmark	90%, 95%, 97.5% and 100%
Number of agents	200
Network Type	Ring
Original number neighbors (links)	4, 400
Lag	3
Runs per setting	60

For ease of reading, Figures 1 through 4 were compiled in the same manner. By way of explanation, the facet columns relate to the rewire settings while the rows relate to the tolerance setting, as per Table 1. Thereby, each individual plot represents a unique combination of these two variables. There are four lines in each plot representing the different values of the discount factor (DF) variable (which is labelled as the % benchmark). The lines in each plot represent the median value of the given variable determined from the 60 runs of the given simulation combination.

Figure 1 illustrates the evolution of the level of trust amongst the population across the simulation. The overarching observation is that trust could not be maintained at the initiation level (1) regardless of which setting is considered. However, a consistent theme is that a higher DF variable (% discount) sees greater depreciation in the trust level—that is, the 100% line marks the lowest trust level. This is explainable by the fact that as the DF variable increases, from say 95% to 97.5%, the standard by which the agents rate their neighbour increases, which in turn increases the rate at each trust is lost.

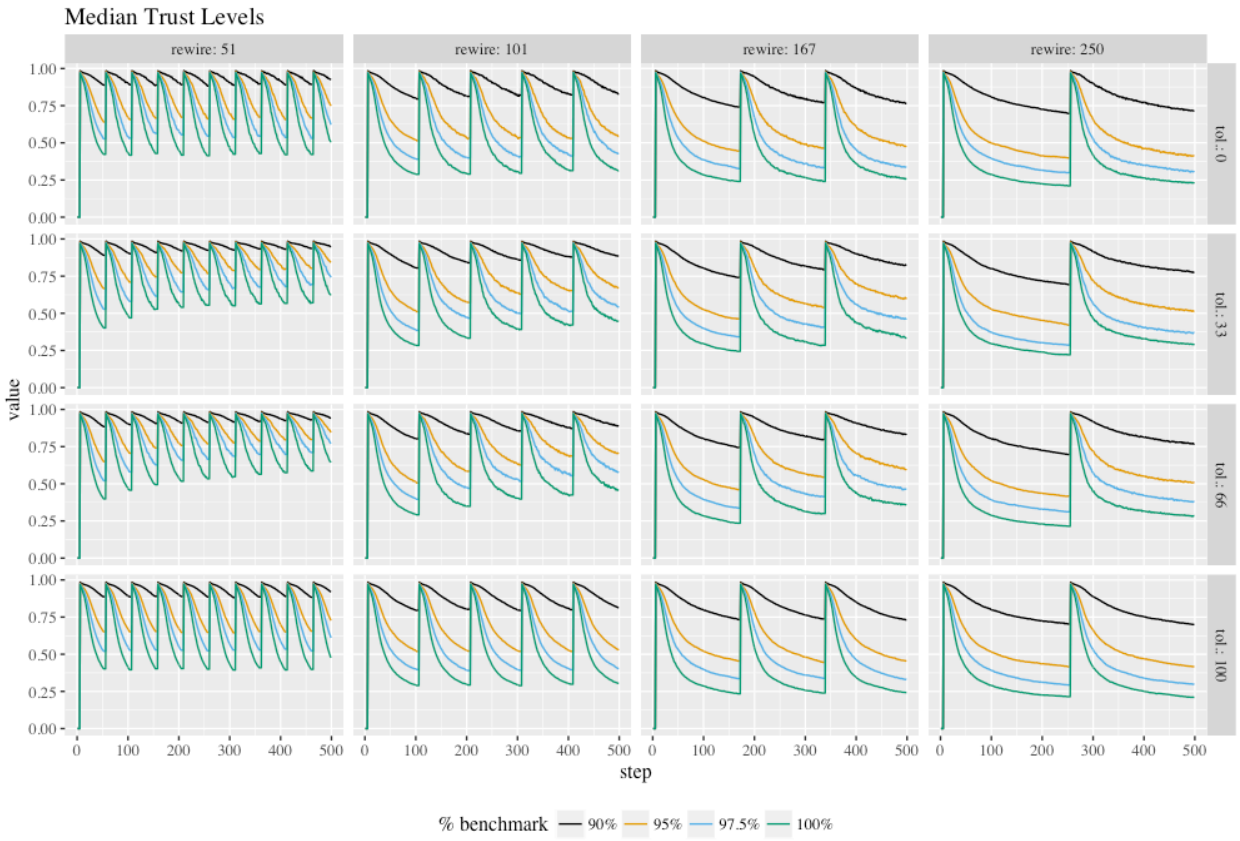

Figure 1: Plots of the median trust level under the various parameter combinations.

In terms of the variation caused by changes in the tolerance and rewiring variables, the following comments can be made. First, the tolerance level appears to have a mixed effect. For the two extreme values of 0% and 100% (total intolerance and total tolerance) they return the poorest performance, and are almost identical. For the two intermediate levels, the level of trust is maintained at a higher level. This result would indicate that by having some but not total tolerance allows you to keep superior neighbours, thus resulting in more trust. This is despite agents resetting their trust in each neighbour to 1, regardless of any history during the rewiring stage. Future iterations of the model could improve upon this by maintaining the trust levels of any existing neighbours. It can also be seen that by increasing the number of switches (see Figure 2) made by each agent does not improve the level of trust, or at least the rate at which it is maintained. For a given level of tolerance, by increasing the interval at which agents select new neighbours has a detrimental effect on the trust level. This a direct function of resetting the trust level.

The purpose of Figure 2 is to illustrates the cumulative number of switches performed by the agents. The rationale for the metric is that it provides a proxy for the willingness of the agents to disregard their existing neighbours and find superior ones. The first obvious observation is that if agents have a 100% tolerance for the information performance of their neighbours, they will not make any switches. In combination with Figure 1, this provides an interesting point, which is that by switching all your neighbours all the time produces similar results to maintaining your neighbours for the long term. Another result of note is that having a higher benchmark by which you judge your neighbours, results in more switches for any given combination of the rewiring and tolerance variables. The intuition is straightforward—you are more likely to disconnect with a neighbour when you hold them to a higher standard.

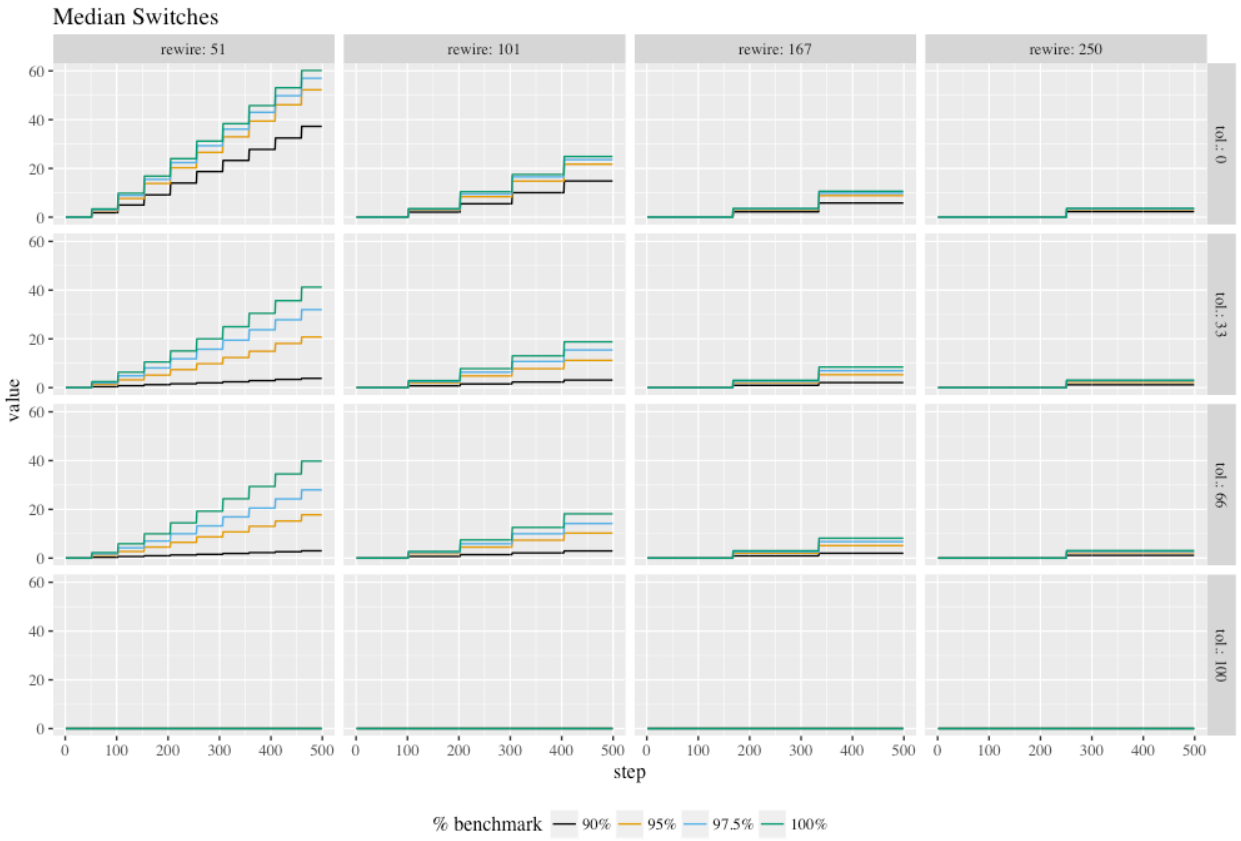

Figure 2: Plots of the median number of switches under the various parameter combinations.

Figure 3 illustrates the number of links that are recorded in the network. The number of links grows as agents disconnect with their existing neighbours and pursue better performing agents. An important point to note is that, given the rewiring process sees agents form un-directed links with

other agents, the number of neighbours that an agent receives information from grows as a function of who the agent chooses to be a neighbour with, as well as of who chooses to be a neighbour with the given agent. This is because an agent automatically receives information from an agent that connects with them. Future iterations of the model could adjust this by using directed links so that an agent only receives information from the agents they specifically decide to connect with.

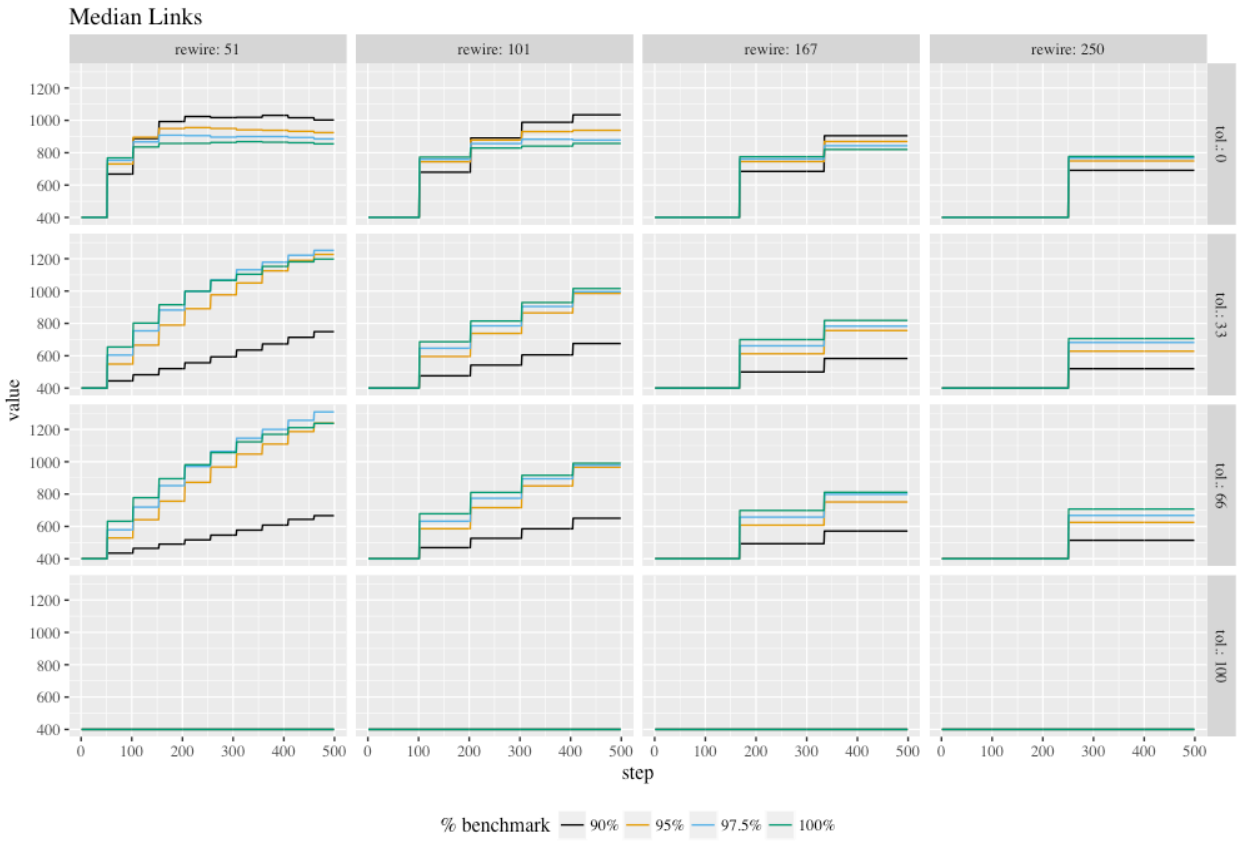

Figure 3: Plots of the median number of links under the various parameter combinations.

In line with the observations from Figure 2, a higher benchmark sees a higher number of links for a given combination of the rewiring and tolerance variables, because of greater switching. The combination of a short rewiring interval (51), and no tolerance (the top left plot) produces an interesting result in that the number of links appears to find a steady state, while all other combinations have a monotonically increasing number of links. This is suggestive of the fact that, in the given environment, there is an optimal number of links that exist. Future work could look to explore this finding.

The final plot illustrates the evolution of the clustering coefficient of the network. The clustering coefficient is an important metric because it measures the degree by which the agents in the network tend to cluster together. High clustering is suggestive of agents forming close knit groups. The initial setting for the creation of the ring network has each agent joined to their four closest neighbours resulted in an averaging clustering coefficient of .5. Therefore, the question is how the various parameter combination affected the coefficient. A common finding across all combinations where agents selected new agents was that at the first rewiring step the clustering coefficient collapses as agents abandon their existing well-structured network in pursuit of the better performing agents. As additional re-wiring steps are undertaken, it is seen that in most cases the coefficient starts to increase. This is indicative of more and more agents connecting to similar high performing agents, and therefore the network evolving a level of structure.

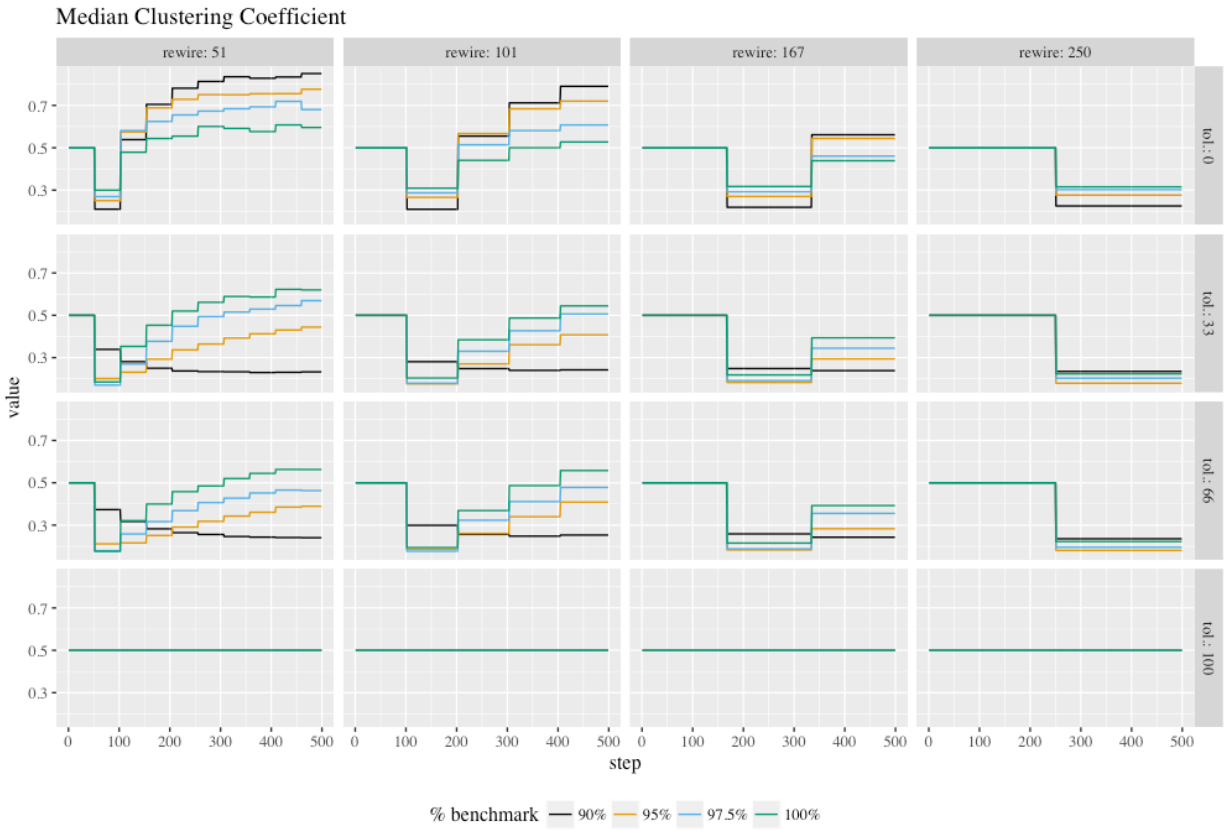

Figure 4: Plots of the median clustering coefficient under the various parameter combinations.

The one interesting exception is where agents exhibit some tolerance for their neighbours, and they hold them to lower standards (a benchmark of 90%), then the clustering coefficient does not

improve. A possible explanation for this comes by combining the observations from Figure 2 and 3, that is the same combination of parameters sees agents make less switches, resulting in less links. At this point any direct implications for the level of trust in the network cannot be drawn.

Extensions and Potential Applications

The above model addresses how information and trust flow across an evolving network. There are several possible extensions, including:

- Rather than initiating the model with a ring network, a small-world, random or scale-free network could be utilized. Then the evolution of the networks could be compared with the intention of establishing which network structure is more stable;
- As mentioned in the results section, undirected links could be replaced with directed links so that agents only received, and assess, information from the agents of their choosing. Again, a comparative analysis of the network evolution would be an interesting stream of research. Additionally, the evolution of the trust in the network would be an area of curiosity;
- The last extension involves not resetting the trust level of existing neighbours at the re-wiring step, with the trust and network evolution consequences would be the main area of interest.

As for real-world applications, the model could be modified to capture any circumstances where the exchange of information is a vital consideration. Key examples would be interactions of market participants who exchange information to decide about their investments or between countries that use information to form their preferences about compliance with international agreements.