EXAMINING THE RELATIVE DISAGREEMENT MODEL OF OPINION DYNAMICS WITH KLEMM-EGUILUZ SOCIAL NETWORK TOPOLOGIES

Michael Meadows

University of Bristol
Department of Computer Science
University of Bristol
The Merchant Venturers Building
Woodland Road
Bristol
BS8 1UB
United Kingdom

michaeljmeadows86@googlemail.com

ABSTRACT
This paper presents a brief history of models of opinion dynamics and summaries of the work from the creation of the Bounded Confidence (BC) model through to the more recent development of the Relative Agreement (RA) model and finally of the Relative Disagreement (RD) model. As a result of the re-examination and correction of the original specification of the RA model given by Meadows and Cliff, and subsequent first investigation of the RA model operating within non-trivial but realistic social networks the RD model was proposed as not only an extension but a significant improvement. Given that these two highly successful approaches have been taken with the RA model, it is now necessary to present a full exploration of the new RD model operating within the same non-trivial topologies.

Keywords: relative disagreement model, opinion dynamics, Klemm-Eguiluz networks, extremist behaviour

1. INTRODUCTION

1.1. Opinion Dynamics
The term "opinion dynamics" has come to cover a broad range of different models applicable to many fields ranging from sociological phenomena to ethology and physics (Lorenz 2007). The focus of this paper is on an improvement to the "Relative Disgreement" model (Meadows and Cliff 2013b), that was originally developed as an improvement to a model designed to assess the dynamics of political, religious and ideological extremist opinions, and the circumstances under which those opinions can rise to dominance via processes of self-organisation (i.e., purely by local interactions among members of a population) rather than via exogenous influence (i.e. where the opinion of each member of a population is influenced directly by an external factor, such as mass-media propaganda). The RA model was developed with the aim of helping to explain and understand the growth of extremism in human populations, an issue of particular societal relevance in recent decades where extremists of various religious or political beliefs have been linked with significant terrorist acts.

Suppose a group of n experts are tasked with reaching an agreement on a given subject. Initially, all the experts will possess an opinion that for simplicity we imagine can be represented as a real number x, marking a point on some continuum. During the course of their meeting, the experts present their opinion to the group in turn and then modify their own opinion in light of the views of the others, by some fixed weight. If all opinions are equal after the interaction, it can be said that a consensus has been reached, otherwise another round is required. It was demonstrated by de Groot (1974) that this simple model would always reach a consensus for any positive weight. Although highly abstract and clearly not particularly realistic, this simple model has become the basis for further analysis and subsequent models (e.g. Chatterjee & Seneta 1977; Friedkin 1999).

Building on the de Groot model, the Bounded Confidence (BC) model included the additional constraint that the experts would only consider the opinions of others that were not too dissimilar from their own (Krause 2000); this is also known as the Hegselmann-Krause model. The BC model adds the idea that each expert has a quantifiable conviction about their opinion, their uncertainty, u. It was demonstrated that although a consensus may be reached in the BC model, it is not guaranteed (Hegselmann & Krause 2002). It was observed that when the BC model is set in motion with every agent having an initially high confidence (low uncertainty) about their own random opinion, the population disaggregates into large numbers of small clusters; and as the uncertainty was increased, so the dynamics of the model tended towards those of the original de Groot model (Krause 2000). Later, the model was tested with the inclusion of "extremist" agents, defined as individuals having extreme value opinions and very low uncertainties. In the presence of extremists it was found that the population could tend towards two main outcomes: central convergence and bipolar convergence (Hegselmann & Krause 2002). In central convergence, typical when uncertainties are low, the majority of the population clustered around the central, "moderate" opinion. In contrast, when uncertainties were initially high, the moderate population would split into two approximately equal groups one of which would tend towards the positive extreme and the other towards the negative: referred to as bipolar convergence.

Although these two phenomena have occurred in real human societies, there is a third well-known phenomenon that the BC model is unable to exhibit: an initially moderate population tending towards a single extreme (and hence known as single extreme convergence).

Shortly after the publication of the BC model, Deffuant, Amblard, Weisbuch, and Faure (2002) reported their exploration of the BC model and proposed an extension of it which they named the Relative Agreement (RA) model (Deffuant et al. 2002). The RA model was intended to be capable of exhibiting single extreme convergence in its dynamics.

There are two main differences between the RA model and the BC model. The first change is that agents are no longer expressing their opinion to the group as a whole followed by a group-wide opinion update. Instead, in the RA model pairs of agents are randomly chosen to interact and update. This is repeated until stable clusters have formed. The second change relates to how agents update their opinions. In the BC model an agent only accepted an opinion if it fell within the bounds of their own uncertainty, and the weight that was applied to that opinion was fixed. In the RA model however, an opinion is weighted proportional to the degree of overlap between the uncertainties of the two interacting agents.

These changes represent a push for increased realism. In large populations, individuals cannot necessarily consider the opinion of every other agent; therefore paired interactions are far more plausible. More importantly, the RA model also allows for agents with strong convictions to be far more convincing than those who are uncertain (Deffuant 2006). Thus, although the RA model is stochastic, the only random element of the model is in the selection of the individuals for the paired interactions (Lorenz 2005). As expected, the RA model was able to almost completely replicate the key results of the BC model (Deffuant *et al*. 2000).

Having demonstrated that RA model was comparable to the BC model under normal circumstances, Deffuant *et al.* then added the extremist agents to the population, to see if they could cause shifts in the opinions of entire population. An extremist was defined as an agent with an opinion above 0.8 or below -0.8 and with a very low uncertainty. Conversely, a moderate agent is one whose absolute opinion value is less than 0.8 and with a fixed, higher uncertainty who is therefore more willing to be persuaded by other agents. Under these circumstances, Deffuant *et al*. reported that there are large areas of parameter space in which all three main types of population convergence could occur. The fact that the RA model offers realistic parameter-settings under which single extreme convergence regularly occurs is a particularly novel attraction.

To classify population convergences, Deffuant et al. (2002) introduced the y metric, defined as: y = p'+2 + p'-2 where p'+ and p'- are the proportion of initially moderate agents that have finished with an opinion that is classified as extreme at the positive and negative end of the scale respectively. Thus, central, bipolar and single extreme convergences have y values of 0.0, 0.5 and 1.0, respectively.

While these findings are particularly striking, it raises the question of where the initially extreme agents may have come from. Interestingly, an answer presents itself from the field of psychology. *Social Judgement Theory* states that the opinions of others may fall within a *latitude of acceptance* in which case we may see a converging opinion update (Sherif and Hovland 1961). Conversely, opinions may fall within the *latitude of rejection*, which may result in a diverging opinion update. It is clear that the models given up to now, only one of these dynamics has been taken into consideration. Thus, the Relative Disagreement (RD) model was created (Meadows and Cliff 2013b). With this model it was shown that by utilising an analogous dynamic for quantifying disagreement as with agreement in the RA model, it was possible to replicate all three population convergences without the artificial need for extremist agents. A full specification of this model is given in the next section, as it is central to the research presented in this paper.

1.2. Social Networks

While much of the work described previously has taken place with agents represented as nodes on a fully connected graph, there is a growing movement towards examining these models under non-trivial topologies. Small World (SW) networks were introduced by Watts & Strogatz (1998), and a full introduction is beyond the scope of this paper. Suffice it to say that SW networks exhibit both low average path lengths and social clustering. Watts and Strogatz introduced an attractively simple stochastic algorithm for constructing SW networks. Nevertheless, one limitation of SW networks as models of human social networks is the extent to which SW networks have unrealistic degree distributions. In real social networks, the majority of nodes often have few connections, while a small number have very high degrees. A well-known possible resolution of this was proposed by Barabási and Albert (1999). The Barabási-Albert (BA) algorithm could construct random graphs with low average path lengths that also obeyed a power law in degree distributions (scale free networks). However, the BA model is unable to generate networks that exhibit clustering levels as high as those in observed social networks and so, although both SW and BA models were useful as research tools, neither could claim to be entirely realistic.

To construct a graph that would exhibit all three qualities observed in real social networks (short average path lengths, high clustering, and a power-law degree-distribution) algorithms have been developed that produce hybrid networks that mix SW and BA characteristics. The KE algorithm introduced by Klemm & Eguíluz (2002) is the one used in this paper. The KE model begins by taking a fully connected graph of size m, the nodes of which are all initially considered active. A network is then "grown" by adding nodes iteratively to all of the currently active nodes in the graph after which a random active node is deactivated and the newest node is assigned to be active. When adding these nodes however, with a probability μ_{KE} each new connection the node forms is assigned to a node using preferential treatment (a node with a higher degree is more likely to be randomly chosen) as in the BA model. With this addition, we see that when μ_{KE}=1.0 the resulting network is identical to the BA model and with μ_{KE}=0.0 the network is generated with topological characteristics as in the SW model. As Klemm and Eguiluz (2002) note, for values of μ_{KE} between 0.0 and 1.0, KE networks exhibit properties that are "hybrid" mixes of the properties of SW and BA networks. For that reason, in this paper we use KE networks to explore the dynamics of the RA model in nontrivially structured populations.

2. SPECIFICATION

For completeness, it is important to now provide a full specification of the RD model to allow for replication and extension work. Returning to the population of *n* agents, each individual *i* is in possession of two variables; an opinion *x*, and an uncertainty *u*, both of which are represented by real numbers. In the RA model, opinion was initially set in the range of -1.0 to 1.0, with extremists being defined as agents whose opinions lay below -0.8 or above 0.8. As the goal of the RD model was to replicate the behaviour seen in the RA model but without extremist agents, opinions may not be initially set outside the range of -0.8 and 0.8, although the maximal values are retained from before. With no extremist agents, there is no longer any constraint on uncertainties used and so they are assigned randomly using a simple method to bias agents towards being uncertain (as it is in uncertain populations that more interesting results are to be found) given by:

$$u = min(random(0.2, 2.0) + random(0.0, 1.0), 2.0)$$

Random paired interactions take place between agents until a stable opinion state is produced. Unlike in the original RA and RD models, this no longer means two agents randomly chosen from the population but instead requires taking one agent at random followed by a randomly chosen neighbour of that agent as defined by the KE network. The relative agreement between agents *i* and *j* is calculated as in the RA model by taking the overlap between the two agents' bounds h_{ij}, given by:

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

Followed by subtracting the size of the non-overlapping part given by:

$$2u_i - h_{ij}$$

So the total agreement between the two agents is given as:

$$h_{ij} - (2u_i - h_{ij}) = 2(h_{ij} - u_i)$$

Once that is calculated, the relative agreement is then given by:

$$2(h_{ij} - u_i) \,/\, 2u_i = (h_{ij} / u_i) - 1$$

Then if $h_{ij} > u_i$, then update of x_j and u_j is given by:

$$x_j := x_{\underline{i}} + \mu_{RA}[(h_{ij} / u_i) - 1](x_i - x_j)$$
$$u_j := u_{\underline{i}} + \mu_{RA}[(h_{ij} / u_i) - 1](u_i - u_j)$$

Similarly, the relative disagreement between agents *i* and *j* is calculated by a very similar method to find g_{ij}:

$$g_{ij} = max(x_i - u_i, x_j - u_j) - min(x_i + u_i, x_j + u_j)$$

Followed by subtracting the size of the non-overlapping part given by:

$$2u_i - g_{ij}$$

So the total disagreement between the two agents is given as:

$$g_{ij} - (2u_i - g_{ij}) = 2(g_{ij} - u_i)$$

Once that is calculated, the relative disagreement is then given by:

$$2(g_{ij} - u_i) \,/\, 2u_i = (g_{ij} / u_i) - 1$$

An analogous method for calculating the agents' disagreement was chosen for ease of understanding as it also facilitates the need for calculating relative disagreement. Given that we would not want the disagreement update to occur in every instance of disagreement, as SJT suggests that this would not occur in every real-world instance of disagreement. Therefore if $g_{ij} > u_i$ and with a probability λ, the update of x_j and u_j is given by:

$$x_j := x_{\underline{i}} - \mu_{RD}[(g_{ij} / u_i) - 1](x_i - x_j)$$
$$u_j := u_{\underline{i}} + \mu_{RD}[(g_{ij} / u_i) - 1](u_i - u_j)$$

Note that if all disagreement updates are capped so that agents' opinions may not exceed the initial bounds of -1.0 to 1.0.

3. PARAMETERS

3.1. Size Variation

Firstly, it is necessary to analyse the variations as the scale of the model is altered. If the work is applicable for all values of n, we are often able to gain the clearest picture with large populations, as the level of noise is proportionally lower. To that end, Figure 1 shows how the overall dynamic of the model is altered as n increases with varying values of μ_{RD} with $\lambda = 1.0$.

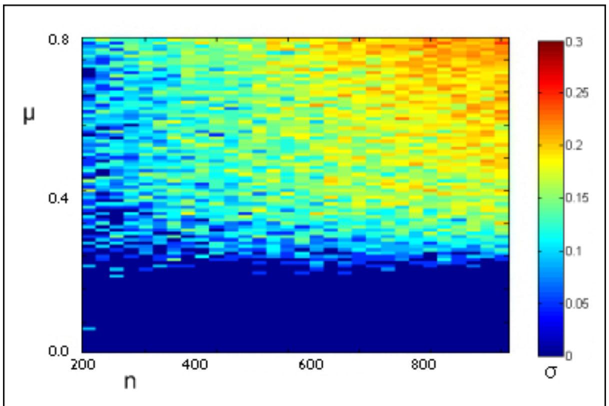

Figure 1: Average standard deviation heat map for varying values of μ_{RA} and μ_{RD} (treated as a single value μ) and n, when agents have randomised values of u, with $\lambda = 1.0$, $m = 6$, and $\mu_{KE} = 0.6$.

Here we see some very interesting, and possibly unexpected results. If we compare with the effect on the RA model (Meadows and Cliff 2013a), it was observed that as n grew, the stability of the population was increased. Much like with the later discussion of μ_{KE}, it can be seen that the influence of extremist agents is weakened because, in order to have an impact, they must rely on the subsequent influence of the moderate agents with which they interact. In the RD model, this is no longer a factor as disagreement interactions can be caused, in theory, by all agents (although in practice some agents will have such large uncertainties coupled with central opinions that they rarely cause a disagreement). Thus we see that the increase in stability typically caused by increasing n in the RA model with KE networks, is missing here.

3.2. Disagreement Probability λ

It is clear that when we set μ_{RD} to 0.0 no disagreement updates are possible, and similarly setting λ to 0.0, causes a similar result as no disagreement can lead to an update. Thus, when comparing extreme values of λ we must use low, but non-zero values, and high values of λ as shown in Figure 2.

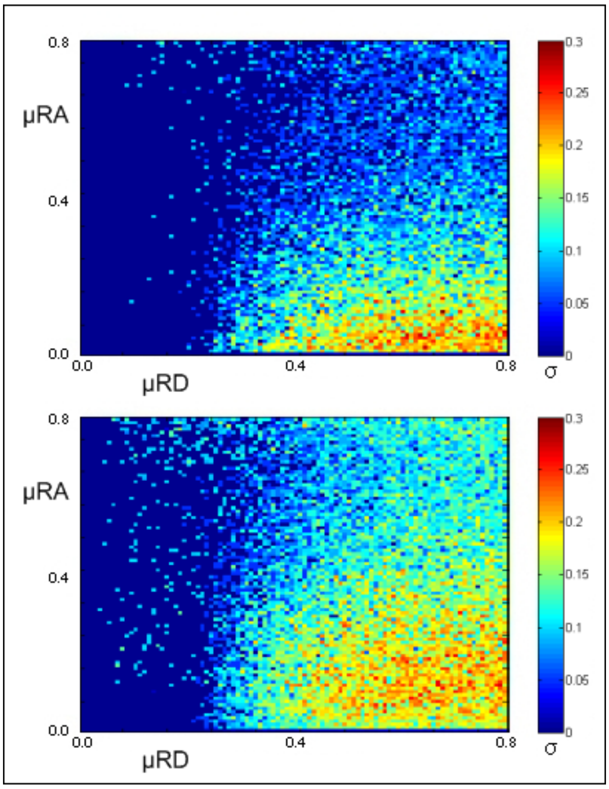

Figure 2: Average standard deviation heat map for varying values of μ_{RA} and μ_{RD} when agents have randomised values of u, $m = 6$, $n = 200$ and $\mu_{KE} = 0.5$. In the top heat map $\lambda = 0.25$ and in the bottom $\lambda = 1.0$.

As one may expect, an increase in λ leads directly to an increase in population instability. Given that when $\lambda = 0.0$ there are no disagreements that cause a repelling update occurring over the course of the simulation, we can understand that with a low value of λ there will similarly be a low proportion of repelling updates. However, with the λ set to its maximal value, every disagreement leads to an update and so we see a greater level of instability in the population. This clearly makes sense. What is surprising is that it is possible to observe the overriding importance of the disagreement update weight μ_{RD}. As before, we can see that when $\mu_{RD} > 0.2$ there is a significant increase in average instability, regardless of λ. This is interesting as it means that the effect of a disagreement must have a minimum impact in order to be influential.

3.3. Initial Network Size m

To further aid the comparison of the RD model with the RA model, consideration must be made to the effect of the initial network size m, from the KE algorithm. Figure 3 presents this analysis of the weight μ (where $\mu = \mu_{RA} = \mu_{RD}$) and m.

influence over the majority of moderates through the moderates with which they are connected).

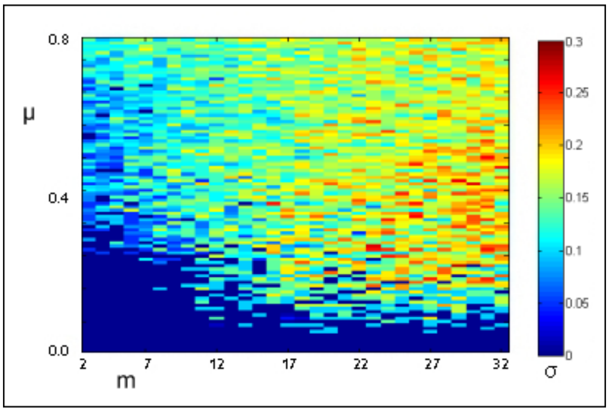

Figure 3: Average standard deviation heat map for varying values of μ_{RA} and μ_{RD} (treated as a single value μ) and m, when agents have randomised values of u, with $\lambda = 1.0$, $n = 200$ and $\mu_{KE} = 0.5$.

The first observation must be that instability grows as m increases. The nature of that growth is certainly interesting however. For very low values of m, in all cases of μ_{KE}, we can see that only when the update weights μ_{RA} and μ_{RD} are significant values can any instability be reliably inferred. As m grows, we see that the minimum values required from μ_{RA} and μ_{RD} for instability decrease. This is in line with the finding in Meadows and Cliff (2013a), which discusses how the typical agent degree is linked with the size of the m. It is the agent degree that we rely on most heavily for instability, as we have found repeated examples that show the greater the connectivity within a population the greater the opportunity for instability.

Another interesting observation is that when m increases, the "height" of the unstable region on the heat map graph shows diversity itself. With high values of μ, we see a smoother area with a lower variance, but as μ decreases, we see the data become noisier and with greater diversity in the possible variances. This unexpected result is quite interesting as it shows an interesting nuance of the RD model. The cause of this result can be explained by the fact that the weight μ was lower. When this occurs it allows for a slower convergence (as each interaction is less influential to an agent's opinion), but with a higher value of m, the agents are exposed to multiple viewpoints and as such can allow for greater swings in the possible convergence.

3.4. Mixing Parameter μ_{KE}

In Meadows and Cliff (2013a) it can be seen most clearly that the RA model operating within a clustered population was far more stable when compared to a BA network and certainly more so than a fully connected graph. Here the main reason was that the RA model's extremist agents struggled to exert their influence over longer path lengths (i.e. they only able to exert their influence over the majority of moderates through the moderates with which they are connected).

In contrast we can see very clearly in Figure 4 that the effect of clustering, although present, is severely limited. When compared to the RA model's behaviour in a clustered population however, this effect is barely noticeable.

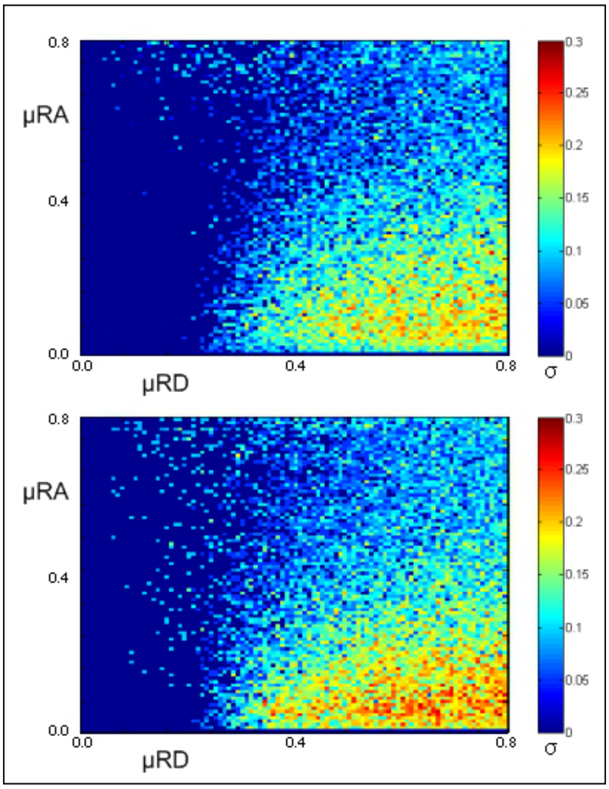

Figure 4: Average standard deviation heat map for varying values of μ_{RA} and μ_{RD} when agents have randomised values of u, $m = 6$, $n = 200$ and $\lambda = 0.5$. In the top heat map $\mu_{KE} = 0.0$ and in the bottom $\mu_{KE} = 1.0$.

The cause of this difference lies in the dynamics of the model. In the RA model, a select few extremist agents are responsible for causing overall instability, which results in their effects being severely limited when their reach is similarly cut off. This does represent an interesting insight into the dynamics of opinion exchange in real life, considering that real world social networks are themselves highly clustered. Given that it has been established that real world terrorist networks exhibit high levels of clustering and that they are a particularly resilient to external influence, the RA model's contribution to study is clearly of merit. However, this outlook fails to explain how in the real world, highly clustered populations can still lead to examples of single extreme or bipolar convergence.

The RD model offers a solution to this difficulty by the way that extremism propagates. Instead of only a limited number of agents causing instability, the instability is a product of the interaction of every agent. That is, any two agents that would otherwise be considered to be moderate can still produce instability through a disagreement. Thus we see that even in clustered populations, when compared with BA networks, agents have approximately equal average degrees, and so the effect of clustering is almost negated. The effect is not entirely ruled out because when agents are pushed towards extremism, their influence on other agents is still dependent on those that they have connections to. Therefore clustering still plays a role in maintaining a slight degree of stability in the population, but simply not to the same level as with the RA model.

While it should not be claimed that this particular difference between the RA and RD models shows an improvement over the RA model, it is an interesting and alternative insight. Instead the two models should be thought of as analysing different aspects of the same problem, with the relevant dynamic being used as required.

4. FURTHER WORK

It is clear that there is much that can be learned from the RD model with the primary focus being to further understanding the dynamics involved in the creation and propagation of extremism in a population. In particular, this application has been most closely linked to the study of terrorist development. The link between this field and the abstract work is clear, however what is currently lacking is a concrete demonstration of the parallels that validate this theoretical work. While the body of work that has been produced is too large to be applied directly to the empirical evidence, it is obvious that researchers of the abstract must find ways to establish the validity of their research. This "application work" has been largely lacking in many fields, including opinion dynamics (Sobkowicz 2009). Therefore, demonstrating real world examples of behaviours exhibited by the RA and RD models, in particular the convergence types, in relation to political and ideological extremism is the most important step for real world validation.

Although it is clear that that the given work is most easily applied to these areas, this arbitrary restriction limits the potential use of the RA and RD models. By expanding the real world applications of these models, it can be seen that many areas are in fact overlapping, although potentially unaware of their related qualities. Ignoring surprisingly related bodies of work because "it is from a different field" or because researchers are simple unaware of each other's work, is not an acceptable justification. What one field has learnt should not need to be relearnt by another. Therefore, finding and demonstrating these overlaps in knowledge is crucial and should be a main focus for future work.

5. CONCLUSION

There were many aims for this work, including demonstrating further reliability of the RD model from its earlier introduction. The most crucial element of this proof relies on the RD model operating under these constraints maintaining comparable behaviour to both the RA model under KE topologies as well as its own behaviour with a fully connected network It has been found that even under non-trivial topologies, the RD model is still capable of producing instances of all three types of convergence. While it is also clear that the disagreement interaction is more influential for instability than the agreement, both are required for single extreme convergence. Replicating the three convergences is clearly the most basic requirement that any model hoping to improve upon the RA model must satisfy. Furthermore, it is possible to see similar overall dynamics in the behaviour of the RD model when compared with its operation in a fully connected graph.

Given that the RA model maintained comparable behaviour over various networks, it is encouraging that the same can be said for the RD model.

In comparison to the RA model itself, it is clear that many of the dynamics that applied to the RA model operating within Klemm-Eguíluz networks apply also to the RD model under the same constraints. We see that in both models when the agent network is not a complete graph, population stability is slightly increased. In addition, when the social network is highly clustered the stability of the population is further increased, however, this further increase is not as great as with the RA model, because of the differences of how extremism is caused in the two models. The RA model shows that a clustered population remains stable and relatively invulnerable to external influence while the RD model shows that influence exerted through disagreements in a social network can still cause extremism.

One of the earlier aims of the RD model was to represent a further push for realism and answer the questions surrounding instability without the need for extremist agents. While it is appears more than reasonable to state that this aim has been fulfilled, it appears that there are a number of subtle differences that can be highlighted between the RA and RD models. Although it would be necessary for differences to be present, the dynamic previously alluded to (that instead of a proportion of agents being the destabilising factor in the RA model, while the whole population causes instability with RD) suggests that there is still work to be done analysing the RA model. The question of how extremism may spontaneously appear without the need for already present extremism is one of the key factors that the RA model has been unable to answer. The RD model offers one possible, and very plausible explanation based on empirical evidence. Also, showing how highly clustered populations may still result in a single extreme convergence, something that clearly has happened in the real world, implies that the improvements of the RD model are worth further research.

It is prudent, therefore, to state that although the RD model offers a number of realistic and useful improvements over the RA model, both models represent useful tools to further learn about the dynamic of opinion spread.

Before any further work should be undertaken it is important to validate the findings and dynamics of both the RA and RD models. As has been discussed in the previous section, there are many disparate fields that are pursuing very similar lines of research apparently oblivious to other work that may be of use. For that reason, highlighting further applicability of these models, as well as looking for further validation of the models themselves, is an essential step in the development of our understanding of these dynamics.

REFERENCES

Barabási, A. and Albert, R., 1999. Emergence of Scaling in Random Networks. *Science*, **286:** 509-512.

Chatterjee, S., and Seneta, E., 1977. Towards Consensus: Some Convergence Theorems on Repeated Averaging. *Journal of Applied Probability*, **14**(1): 88-97.

De Groot, M., 1974. Reaching a Consensus. *Journal of the American Statistical Association*, **69**(345): 118-121.

Deffuant, G., Amblard, F., Weisbuch, G., and Faure, T., 2002. How can extremism prevail? A study based on the relative agreement interaction model. *Journal of Artificial Societies and Social Simulation*. 5(4):1. Available from: http://jasss.soc.surrey.ac.uk/5/4/1.html

Deffuant, G., 2006. Comparing Extremism Propagation Patterns in Continuous Opinion Models. *Journal of Artificial Societies and Social Simulation*, 9(3):8. Available from: http://jasss.soc.surrey.ac.uk/9/3/8.html.

Friedkin, N., 1999. Choice Shift and Group Polarization. *American Sociological Review*, **64**(6): 856-875.

Hegselmann, R., and Krause, U., 2002. Opinion dynamics and bounded confidence: models, analysis and simulation. *Journal of Artificial Societies and Social Simulation,* 5(3):2. Available from: http://jasss.soc.surrey.ac.uk/5/3/2.html

Klemm, K., and Eguíluz, V., 2002. Growing Scale-Free Networks with Small-World Behavior. *Physical Review E,* **65**: 057102.

Krause, U., 2000. A Discrete Nonlinear and Non-Autonomous Model of Consensus Formation. Communications in difference equations: *Proceedings of the Fourth International Conference on Difference Equations*, 227-236, August 27-31, Poznan, Poland.

Lorenz, J., and Deffuant, G., 2005. The role of network topology on extremism propagation with the relative agreement opinion dynamics. *Physica A: Statistical Mechanics and its Applications* **343**: 725-738.

Lorenz, J., 2007. Continuous Opinion Dynamics under Bounded Confidence: A Survey. *International Journal of Modern Physics C* 18:1-20.

Meadows, M., and Cliff, D., 2012. Reexamining the Relative Agreement Model of Opinion Dynamics. *Journal of Artificial Societies and Social Simulation,* 15(4):4. Available from: http://jasss.soc.surrey.ac.uk/15/4/4.html

Meadows, M. and Cliff, D., 2013a. The Relative Agreement Model of Opinion Dynamics in Populations with Complex Social Network Structure. *Complex Networks IV,* 71-79, March 13-15, Berlin, Germany.

Meadows, M. and Cliff, D., 2013b. The Relative Disagreement model of opinion dynamics: where do extremists come from? *7th International Workshop on Self-Organizing Systems,* May 9-10, Palma de Mallorca, Spain.

Sherif, M., and Hovland, C., 1961. *Social Judgement: Assimilation and contrast in communication and attitude change*. New Haven, Connecticut: Yale University Press.

Sobkowicz, P., 2009. Modelling Opinion Formation with Physics Tools: Call for Closer Link with Reality. *Journal of Artificial Societies and Social Simulation,* 12(1):11. Available from: http://jasss.soc.surrey.ac.uk/12/1/11.html

Tsvetovat, M., and Carley, K., 2002. Knowing the Enemy: A Simulation of Terrorist Organisations and Counter-Terrorism Strategies. *Proceedings of CASOS Conference 2002, Day 3*, Electronic Publication, Pittsburgh, PA.

Watts, D., and Strogatz, S., 1998. Collective Dynamics of Small-World Networks. *Nature,* **393**: 440-442.

AUTHORS BIOGRAPHY

Michael Meadows is a PhD Student at the University of Bristol. He achieved his BSc in Computer Science at the University of Leicester and then moved to Bristol to pursue research more closely linked to his own interests. After starting on a joint venture research programme with the University of Bristol and BAE Systems, Michael joined the LSCITS group and is currently supervised by Professor Dave Cliff. He is currently researching opinion dynamics across large populations from the broader subject of examining large-scale attacks on socio-technical systems.