

## **UC Merced**

### **Proceedings of the Annual Meeting of the Cognitive Science Society**

#### **Title**

Opinion Cascades and Echo-Chambers in Online Networks:  
A Proof of Concept Agent-Based Model

#### **Permalink**

<https://escholarship.org/uc/item/4m97p67t>

#### **Journal**

Proceedings of the Annual Meeting of the Cognitive Science Society, 39(0)

#### **Author**

Pilditch, Toby D.

#### **Publication Date**

2017

Peer reviewed

# Opinion Cascades and Echo-Chambers in Online Networks: A Proof of Concept Agent-Based Model

Toby D. Pilditch (t.pilditch@ucl.ac.uk)

Department of Experimental Psychology, University College London  
WC1E 6BT, Gower Street, London, United Kingdom

## Abstract

In online networks, the polarization of opinions (e.g., regarding presidential elections or referenda) has been associated with the creation of “echo-chambers” of like-minded peers, secluded from those of contrary viewpoints. Previous work has commonly attributed such phenomena to self-regarding preferences (e.g., confirmation bias), individual differences, and the pre-dispositions of users, with clusters forming over repeated interactions.

The present work provides a proof of concept Agent-Based Model that demonstrates online networks are susceptible to echo-chambers from a single opinion cascade, due to the spatiotemporal order induced by lateral transmission. This susceptibility is found to vary as a function of degree of interconnectivity and opinion strength. Critically, such effects are found despite globally proportionate levels of opinions, equally rational agents (i.e. absent conformity, confirmation bias or pre-disposition architecture), and prior to cyclical interactions.

The assumptions and implications of this work, including the value of Agent-Based Modelling to cognitive psychology, are discussed.

**Keywords:** Information cascades; opinion dynamics; belief updating; Agent-Based Models

## Introduction

As online networks, such as social media, have developed and increased in popularity, research regarding the spread of false information, the polarization of opinions (Dandekar, Goel, & Lee, 2013), and echo-chamber phenomena (Del et al., 2016) have become increasingly pertinent topics. Such phenomena pose a problem to society, and democracy as a whole, given the average user’s exposure to only the information and opinions of their *local* (i.e. direct) network, leading to a break-down in informed debate and consensus.

Recently, questions regarding how individuals on a network receive new information and form or adopt opinions has come to the fore. Whether on topics of national referenda, deciding between presidential candidates, or interpreting news events (e.g., who is at fault in the annexation of Crimea, the shooting down of passenger aircraft, the political correctness of a reported comment or behavior), it has become more and more common for such information to be ascertained via social media<sup>1</sup>. In this way, an agent’s source of information comes through a filter of network-acquaintances, presenting an unprecedented degree

of lateral, peer-to-peer dissemination of information. Such peer-to-peer transference of information, in a time where the information itself (whether “fake news”, political memes, or posted opinion) carries a bias in its view of the world, presents a problem that psychology and multi-agent modelling is well-placed to answer.

The purpose of the present paper is two-fold: Firstly, this work provides a novel demonstration of the dangers of lateral propagation of opinions in online networks, based solely on the level of interconnectivity and the inherent way in which interpreted events (i.e. opinions) travel through them. This results in high levels of false consensus and echo-chambers on a local level within the network. Critically, such localized clustering is shown to occur *before* any repeated interaction behaviors, and is robust to both different opinion strengths and propensities to communicate. Secondly, this work presents an argument that cognitive science is readily placed to lend insight into these interactive, societal level phenomena, and the super-aggregate behaviors. Such insight can be lent by the ready application of cognitive models taken from individual-based empirical work and theory, to multi-agent simulations, known as Agent-Based Models (ABMs), so as to uncover and explain phenomena beyond the scope of individual-based experiments.

## Cascades and Opinions

How information is communicated between individuals on a societal (or multi-agent) scale, and its consequences, has been formally approached in two main areas; information cascades and opinion dynamics.

Research in information cascades has focused on the way in which information is spread through a system. This has included how networks may be resistant to cascading influence, such as the spread of cultural fads (Watts, 2002). Such work has typically characterized “information” as a singular, memetic entity that is propagated or hindered by either the properties of individuals within the network (such as the proportion of “easily influenced individuals”, see Watts & Dodds, 2007), or the structure of the network itself (e.g., hierarchical influencers; see Watts, 2002). This work has illustrated power law effects in information propagation across networks, an effect akin to percolation theory in physics (for a review, see Essam, 1980), wherein the clustered structure of a system leads to a critical singularity event (i.e. cascade). These cascades result in cluster size distribution effects, where smaller, more numerous clusters

<sup>1</sup> In 2016 a PewResearch poll found the majority (62%) of US adults get their news through social media. Source: <http://www.journalism.org/2016/05/26/news-use-across-social-media-platforms-2016/>

occur as systems become more interconnected (Amar & Family, 1995; Meakin, Vicsek, & Family, 1985).

Research in opinion dynamics has instead focused on the cyclical interactions of individuals within a network. In particular, it has looked at the ways in which individuals and groups interact so as to either reach a consensus (Acemoglu & Ozdaglar, 2011; Hegselmann & Krause, 2002) or segregate into polarized sub-groups of homogenous opinion-holders (Dandekar et al., 2013; Duggins, 2016; Zanette & Gil, 2006). Critically, this research has focused on groups of pre-existing opinion-holders. This work has yielded insights into belief-updating via repeated interaction (such as through the use of the Bounded Confidence Model; Deffuant, Neau, Amblard, & Weisbuch, 2000), along with psychologically based models of behaviors including network pruning (Ngampruetikorn & Stephens, 2015), which provides a plausible pruning mechanism of network contacts, based on a confirmation bias (self-regarding) principle.

The present paper interweaves elements of these two lines of literature, in conjunction with cognitive architecture brought forth from models of learning and communication in cognitive psychology. In particular, agents are encoded with three pieces of cognitive machinery: attention (detecting the public declarations of others); learning (incorporating a communication into a belief-state, and evaluating it against evidence); and decision-making (each choosing whether to make their opinion public based on a decision rule). In this way, all agents within the network are equally rational.

By focusing on universal cognitive architecture on the part of agents (and instead introducing stochasticity to the evidence against which an opinion is evaluated), this work argues that echo-chambers may result solely from the way in which networks are structured, and the spatiotemporal order of lateral opinion transference (i.e. an *opinion cascade*).

The semi-random way in which networks are structured (my relational position to the global network is random, but my method of forming my proximal (direct) connections is rule-based (those whom I know)), runs parallel to work on “small-worlds” (Watts & Strogatz, 1998), which have shown susceptibility to cascades and synchronizability. As such, echo-chambers may occur without reliance on repeated interaction (Acemoglu & Ozdaglar, 2011; Duggins, 2016), or individual differences encoded in agents, such as differences in susceptibility, or pre-dispositions towards an opinion (Watts & Dodds, 2007) or hierarchy (see Quattrociochi, Caldarelli, & Scala, 2014).

### Agent-Based Modeling

ABMs are multi-agent, dynamic simulations which use combinations of three central components; agents, patches, and links. *Agents* are the individual actors within a model, and in the present paper, represent individuals within a network. Agents may be encoded cognitive rules (e.g., learning models), simple behaviors (e.g., signaling to

neighboring agents, movement), and values (e.g., prior beliefs, physical positioning). Agents are ascribed various forms of heterogeneity (such as occupying different positions within a network), as multiple agents are generated within the system. As the simulation runs, agents enact behaviors and update their values according to the specific rules ascribed to them, interacting with other agents and the environment accordingly.

Similarly, both *links*, which represent connections between agents, and *patches*, which represent the environment, may be encoded with behaviors and values, and the capacity to dynamically interact and update as the simulation runs. In the present paper, links are used to represent the connections between individuals within a network, and are thus used for signaling between agents. Given the network representation (requiring only agents and the links between them), the present model does not require the use of patches.

ABMs have been used to explore and assess how behaviors on an individual level, when placed within a dynamic, multi-agent, heterogeneous system, can lead to societal level, super-aggregate behaviors (Epstein, 1999, 2006; Schelling, 2006). For example, by encoding a preference in individuals to be neighbors with others who are similar (whether, on racial, socio-economic, or cultural lines), and assuming some stochasticity in signaling such similarity, Thomas Schelling (1971) was able to show the evolution of segregation on a community, and even city-wide level. In a similar manner, the previously mentioned research on information cascades and opinion dynamics (Duggins, 2016) has used this technique to demonstrate a number of phenomena, with relatively few assumptions, that are difficult with traditional, equation-based cognitive modelling.

### A Model of Opinion Cascades

The aim of the current model is to provide a proof that the inherent structure of an online network is susceptible to high degrees of opinion segregation (i.e. false consensuses or echo-chambers). Critically, this segregation does not require repeated interaction, and can instead occur as a consequence of a single “cascade” across a network of rational agents (i.e. assuming no individual differences in cognitive architecture), despite equal proportions of opinion-holders on a global level.

A network of agents is created whereby agents are randomly assigned an XY coordinate, and each outfitted with the cognitive architecture and values outlined below. Each agent then forms links with its neighbors based on proximity in terms of Euclidean distance – representative of relational proximity in online networks (see Duggins, 2016). The number of links agents form is manipulated, and based on the percentage of the total number of agents in the system, from .5%, to 50%. This is calculated by dividing the number of links per agent by the total number of agents in the network. Thus, given a population of 1000 agents, for an interconnectivity of .5%, all agents form links with their

nearest 5 neighbors; for 10%, the nearest 100 agents, and so on. Accordingly, given a fixed population size across simulations, interconnectivity is manipulated via the number of links each agent possesses. In a similar manner, a neutral “event” node is placed in the geographical center of the simulation, and connected to the nearest agents according to the above rules for interconnectivity. Thus, increasing interconnectivity beyond 50% serves no purpose, given that every other agent will have been exposed to the neutral event (i.e. is 1<sup>st</sup> generation), and thus no cascade can occur beyond two time points. Similarly, in the current model, interconnectivity below .5% (i.e. 5 links per Agent) starts to risk fracturing the network into separate entities.

### Cognitive Architecture

Each agent is outfitted with simple cognitive architecture that can be classified into three branches: attention, learning, and propagation.

All agents within the network attend to their linked-neighbors, in that they are sensitive to the first of their neighbors to “declare” an opinion. Having seen such a declaration, the agent then moves into a learning phase to evaluate it.

The communicated opinion thus forms a prior for the evaluating agent. As mentioned previously, the opinions in the model are categorized into a binary division (Opinion A, Opinion B). Thus, from a neutral prior (.5), moving towards Opinion A is assigned a positive direction, whilst moving towards Opinion B a negative direction. In this way, a prior indicating Opinion A should shift the neutral recipient agent positively (e.g.,  $0.5 + 0.1 = 0.6$ ), and negatively for Opinion B. The strength of this shift is accordingly manipulated as a proxy of opinion strength / influence.

To represent the relationship between the strength of an opinion and the likelihood of a recipient adhering to that opinion, a learning model is used that allows agents to evaluate the opinion against stochastic evidence.

Specifically, a reinforcement learning model is used (Rescorla & Wagner, 1972), in which agents evaluate an opinion in light of new evidence, such that the prediction error ( $\delta$ ), multiplied by the learning parameter ( $\beta$ ), is added to the value associated with the opinion (prior) for the current trial ( $Q(t)$ ), leading to an updated opinion value ( $Q(t + 1)$ ).

$$Q(t + 1) = Q(t) + \beta\delta(t) \quad (1)$$

Such models have been adapted (with added complexity) successfully to model the impact of instruction in reinforcement learning (Doll, Jacobs, Sanfey, & Frank, 2009; Staudinger & Büchel, 2013) and are thus considered a suitable placeholder for the proof of concept model. To evaluate the belief, agents then experience a number of evidence trials (arbitrarily set to 10), where evidence values are binary {0, 1}, and are drawn with equal likelihood (i.e.  $P(E=1) = .5$ ). To reiterate, the learning process herein serves as a *representation* for the relationship between prior strength, and its likelihood of acceptance/rejection. Thus, if

the communicated opinion is represented by a weaker prior, it is more likely to be rejected by the learning / evaluation process. Similarly, increasing the amount of available evidence has the equivalent effect of converging the agent to the .5 (neutral) true state of the event (i.e. reducing the likelihood of passing on the original opinion). In this way, stronger opinions make the cascade more deterministic. Further, using a stochastic sampling *process* to dictate opinion uptake serves as a useful baseline model, to which complexity may be added directly to learning processes.

Having evaluated, agents declare for one of the two opinions, based on a decision rule: if  $Q(\text{posterior}) > .5$ , hold Opinion A; if  $< .5$ , hold Opinion B. This declaration is then made public (and thus may act as a prior to attending linked-neighbors) with a probability that is manipulated between systems. For example, a  $P(\text{Declaration})$  of 1 means all agents will make their opinions public, whilst a  $P(\text{Declaration})$  of .1 means there is a 10% probability of agents making their opinion public. This  $P(\text{Declaration})$  bears a parallel to Watts and Dodds (2007) “individual threshold”, found to impact spreading phenomena.

### Dynamics

Given the above architecture has been established, simulations commence with the initial, neutral “event” being witnessed by a portion of the network (based on manipulated interconnectivity). These agents (termed 1<sup>st</sup> generation) start with a neutral prior, and so, based on the stochastic nature of the evidence, half should arrive at each opinion post-evaluation. From this point, if an agent of the 1<sup>st</sup> generation makes their opinion public (based on manipulated  $P(\text{Declaration})$ ), their attentive (presently neutral) linked-neighbors (2<sup>nd</sup> generation) then take this opinion as a prior, and evaluate it according to the procedure above. This 2<sup>nd</sup> generation agents, having come to a decision, then similarly each choose whether to make their opinion public (based on  $P(\text{Declaration})$ ), and thus the 3<sup>rd</sup> generation is exposed. This process continues until there has been no change in the number opinion-holders (of either type) for two consecutive time periods (i.e. if no one has made an opinion public, and thus the opinions have “died out”, or if the network is now completely saturated).

Importantly, for the proof of concept model, having decided upon an opinion, an agent is no longer attentive to further information. This is purposeful to prevent cyclical effects beyond an initial cascade, as the goal of the present paper is to show the susceptibility of interconnected *neutral* agents to an opinion cascade, without resorting to explanations of homophily (Dandekar et al., 2013) and localized consensus reaching (Ngampruetikorn & Stephens, 2015).

For the purpose of the present paper, the behaviors of interest are constrained to two, related measures. Firstly, the global proportion of opinions across the system (i.e. the proportion of agents with Opinion A, and the proportion with Opinion B) is of interest before inferring anything about localized clustering. For example, whether localized

clustering is simply a by-product of a dominant, network-wide opinion. This leads to the second measure: the average percentage of likeminded neighbors an agent possesses. In other words, of an agent's *visible network*, what percentage are in agreement with the agent (e.g., 50% indicates agents directly linked to equal proportions of each opinion-type). The manipulated variables are summarized in table 1 below:


Table 1. System variables

Variable	Description	Levels
Interconnectivity (%)	(Links per Agent / Total Agents in Network) * 100	0.5, 1, 1.5, ... 50
Opinion Strength	Added to (or subtracted from) neutral agent prior (P(H) = .5)	0, 0.1, 0.2
P(Declaration)	Probability of making opinion public	0.1, 0.5, 1

### Central Findings

The above model was implemented in NetLogo (5.2.1). Each system specification (Interconnectivity (100) x Opinion Strength (3) x P(Declaration) (3)) was run independently 100 times, taking an average set of values for each specification. The total number of agents in each simulation was set to 1000.

Figs. 1a & 1b show example outcomes of opinion cascades (A in red, B in blue) across a sparsely connected (1% interconnectivity) and a more densely connected (10% interconnectivity) system, respectively.



Figures 1a and 1b: Sparsely and densely connected networks, post cascade (grey represents unused links).

Importantly, as Fig. 2 illustrates, irrespective of opinion strength, P(Declaration), or interconnectivity, the global proportion of different opinion holders consistently approximates 50/50.




Figure 2: Proportion of opinion holders across network




Figure 3: Degree of Clustering. Calculated as the average percentage of like-minded neighbors an agent possesses (panels represent P(Declaration) conditions).

The degree of clustering (Fig. 3) can be broken down into several key findings. First, and central to the present paper, localized clustering increases as a function of decreasing interconnectivity and opinion strength, with stronger opinions and low interconnectivity (<1%) resulting in the local (directly visible) networks of agents consisting of >90% likeminded individuals<sup>2</sup>. Second, this effect occurs irrespective of the propensity for individuals within the network to make their opinions public<sup>3</sup>. In other words, whether P(Declaration) is at 100% or 10%, localized clustering occurs regardless.

Finally, localized clustering is mitigated by the degree of stochasticity (i.e. as opinion strength moves towards neutral, thus having no communicative impact) and increasing interconnectivity. However, it is important to note that to prevent local clustering requires either no opinion impact or moving towards high (and arguably unrealistic) levels of interconnectivity.

### Discussion

The central finding of the present paper is that the fundamental way in which networks are constructed, when combined with the temporal dynamics of how information travels through them, and the cognitive representation of opinions as a prior, *inherently* leads to false consensus effects and echo-chambers. Thus, the more deterministic peer-to-peer communications are (i.e. how likely is a recipient to adopt the opinion of a sender), and the lower the *relative density* of connections within the network, the greater the impact of the spatiotemporal order (i.e. the larger the cascade sequence) on clustering.<sup>4</sup>

<sup>2</sup> Further simulations in which total network size has been varied, but density has been kept constant at 1% (i.e. 10 agent-links for 1000 agents, 50 links for 5000 agents) have shown clustering effects remain constant (i.e. depend on relational, not absolute links / network size).

<sup>3</sup> Mathematically, P(Declaration) starts to have an impact when it effectively reduces the average number of “functional” links to a point below the absolute threshold for a singular cohesive network (i.e. if it reduces the average number of active links below 4 in the present model; left-hand panel of Fig. 3).

<sup>4</sup> The present model demonstrates this with fixed, neutral (0.5) priors for all agents. If variance in priors is included, such that *SD*

Critically, this effect occurs *prior* to any repeated interactions between agents, separating the present work from opinion dynamic literatures (Acemoglu & Ozdaglar, 2011; Allahverdyan & Galstyan, 2014), and without assuming individual differences on the part of agents (e.g., differences in susceptibility) or singular information types, common to information cascade literatures (Watts, 2002). Further, work in these areas including social network pruning (Ngampruetikorn & Stephens, 2015) and polarization effects (Dandekar et al., 2013; Duggins, 2016), when looking at cyclical interactions, illustrate that repeated interaction is likely to only exacerbate the already high levels of localized clustering.

### False Consensus and Echo-chambers

The effects described in the present work are found to be broadly independent of the propensity to communicate, and robust across the degree of interconnectivity (requiring approximately 50% interconnectivity density to negate, something unfeasible in online networks approaching billions). Putting this into concrete terms, Facebook has an estimated 1.79 billion active users<sup>5</sup>. The average (median) number of “friends” or links is approximately 200<sup>6</sup>, meaning the average user is connected to .000011% of the overall network. To fully negate the effects demonstrated here would require either the severance of lateral transmissions (or decreasing the deterministic capacity of communications sufficiently), or having each user share direct connections with approximately 900 million other users.

The formation of echo-chambers and the polarization of opinions is typically attributed to repeated interaction with a self-regarding preference (Ngampruetikorn & Stephens, 2015) or a signaling of like-mindedness (e.g., trust; see Li, Scaglione, Swami, & Zhao, 2013). This work instead shows that the structure of the network, and the way in which opinions emanate across it, are sufficient to result in false consensus effects and echo-chambers. To put this in more pragmatic terms; regardless of who you know, why you know them, or how you have repeatedly interacted / pruned your network, the fact that you do not, and arguably cannot know *enough* people, no matter who they are, is sufficient to leave you highly susceptible to echo-chambers.

It should be noted that this proof of concept model carries with it several assumptions. Most notably, opinions are classified in a binary fashion, so as to replicate the target

opinion types under investigation, and associated with echo-chambers (e.g., referenda, or political campaigns). Future work is proposed to incorporate variance as they move across a network (i.e. do they dissipate, or become stronger). Secondly, agents attend and evaluate based on the first exposure to an opinion in their immediate network (i.e. those they are directly connected to). Although future work is suggested to incorporate the influence of multiple sources (e.g., via social conformity), such architecture is initially precluded to avoid “baking in” localized clustering effects.

Finally, the present model assumes a flat hierarchy of individuals. Although the argument can be made that fixing the level of interconnectivity for all individuals in a network is an artificial constraint, in terms of the degree of interconnectivity in target systems (e.g., Facebook) the functional difference in interconnectivity among users is between approximately .000011% (200 friends) and .00028% (5000 friends; Facebook user limit). Although structural hierarchy, such as media influencers, may have an impact on dissemination (along with their own motives, such as following pre-existing opinion trends; see Quattrociochi et al., 2014), the present work serves to illustrate that localized clustering can result from the spatiotemporal order of lateral transmission across a network.

### Further Work

The present work, in serving as a proof of concept for an increasingly important phenomenon, and providing some initial assumptions to illustrate the effects in a straightforward manner, leaves the door open for further, more psychologically informed modelling opportunities.

Further work should start to incorporate additional complexity on the part of agent (cognitive) architecture, such as the inclusion of social conformity (Latané, 1981), which is predicted to increase clustering tendencies (and feasibly increase the strength of opinions as they spread throughout the system. Similarly, work on confirmation bias suggests a similarly exacerbating role (Allahverdyan & Galstyan, 2014; Doll et al., 2009; Nickerson, 1998; Staudinger & Büchel, 2013). Finally, the inclusion of Bayesian models of source credibility (Harris & Hahn, 2009; Harris, Hahn, Madsen, & Hsu, 2015; Madsen, 2016) are of interest (Bayesian models of social learning have already started being applied to opinion dynamics; see Acemoglu & Ozdaglar, 2011), given the way in which people form networks (i.e. we tend to select those we know / trust / like when forming our “direct” network). These suggestions are by no means exhaustive, but serve as examples of the promising (and readily applied) further additions to the framework laid out in the present work.

The present work purposefully precludes such psychological elements, which are expected to exacerbate the effects illustrated in this proof of concept model. This choice was made both for reasons of parsimony, and to provide a demonstration that the effects herein do not rely on such processes or explanations.

---

> opinion strength, then clustering severity is reduced. However, this relies on the *strong assumption* that there is independence of opinions across a self-selecting network. If one incorporates instead a degree of dependence in neighbouring opinion-holders, then one has in effect shifted echo-chamber formation to precede opinion transmission, and have thus “baked-in” the result.

<sup>5</sup> Figure taken from monthly active users as of the 3<sup>rd</sup> quarter of 2016. Source: <https://www.statista.com/statistics/264810/number-of-monthly-active-facebook-users-worldwide/>

<sup>6</sup> Figure taken from Pew Research Center survey of Facebook users in 2014. Source: <http://www.pewresearch.org/fact-tank/2014/02/03/6-new-facts-about-facebook/>

In conclusion, the present paper demonstrates that rational agents (i.e. absent special functionality of cognition or individual differences), through the way in which online networks are structured, are intrinsically susceptible to high levels of localized clustering (i.e. echo-chambers) when opinions are transmitted laterally. Further, it is hoped that the present paper serves as an example of how psychological principles taken from the individual level may be applied to a societal level through the use of Agent-Based Models.

### Acknowledgments

This paper has benefitted from comments and discussions with Christos Bechlivanidis and Jens Koed Madsen.

### References

- Acemoglu, D., & Ozdaglar, A. (2011). Opinion Dynamics and Learning in Social Networks. *Dynamic Games Applications, 1*, 3–49.
- Allahverdyan, A. E., & Galstyan, A. (2014). Opinion dynamics with confirmation bias. *PloS One, 9*(7), e99557.
- Amar, J. G., & Family, F. (1995). Critical Cluster Size : Island Morphology and Size Distribution in Submonolayer Epitaxial Growth. *Physics Review Letters, 74*, 2066–2069.
- Dandekar, P., Goel, A., & Lee, D. T. (2013). Biased assimilation, homophily, and the dynamics of polarization. *Proceedings of the National Academy of Sciences of the United States of America, 110*(15), 5791–6.
- Deffuant, G., Neau, D., Amblard, F., & Weisbuch, G. (2000). Mixing beliefs among interacting agents. *Advances in Complex Systems, 3*(01n04), 87–98.
- Del, M., Bessi, A., Zollo, F., Petroni, F., Scala, A., Caldarelli, G., ... Quattrocchi, W. (2016). The spreading of misinformation online. *PNAS, 113*(3), 554–559.
- Doll, B. B., Jacobs, W. J., Sanfey, A. G., & Frank, M. J. (2009). Instructional control of reinforcement learning: A behavioral and neurocomputational investigation. *Brain Research, 1299*, 74–94.
- Duggins, P. (2016). A Psychologically-Motivated Model of Opinion Change with Applications to American Politics. *ArXiv*.
- Epstein, J. M. (1999). Agent-based computational models and generative social science. *Complexity, 4*(5), 41–60.
- Epstein, J. M. (2006). *Generative social science: Studies in agent-based computational modelling*. Princeton University Press.
- Essam, J. W. (1980). Percolation theory. *Reports on Progress in Physics, 43*, 833–912.
- Harris, A. J. L., & Hahn, U. (2009). Bayesian rationality in evaluating multiple testimonies: incorporating the role of coherence. *Journal of Experimental Psychology. Learning, Memory, and Cognition, 35*(5), 1366–1373.
- Harris, A. J. L., Hahn, U., Madsen, J. K., & Hsu, A. S. (2015). The Appeal to Expert Opinion: Quantitative Support for a Bayesian Network Approach. *Cognitive Science, 39*(7), 1–38.
- Hegselmann, R., & Krause, U. (2002). Opinion Dynamics and Bounded Confidence. *Simulation, 5*(3), 2.
- Latané, B. (1981). The psychology of social impact. *American Psychologist, 36*(4), 343–356.
- Li, L., Scaglione, A., Swami, A., & Zhao, Q. (2013). Consensus , Polarization and Clustering of Opinions in Social Networks. *IEEE Journal on Selected Areas in Communications, 31*(6), 1072–1083.
- Madsen, J. K. (2016). Trump supported it?! A Bayesian source credibility model applied to appeals to specific American presidential candidates’ opinions. In *Proceedings of the 38th Annual Meeting of the Cognitive Science Society* (pp. 165–170).
- Meakin, P., Vicsek, T., & Family, F. (1985). Dynamic cluster-size distribution in cluster-cluster aggregation: Effects of cluster diffusivity. *Physical Review, 31*(1), 564–569.
- Ngampruetikorn, V., & Stephens, G. J. (2015). Bias, Belief and Consensus: Collective opinion formation on fluctuating networks. *arXiv Preprint arXiv:1512.09074*.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology, 2*, 175–220.
- Quattrocchi, W., Caldarelli, G., & Scala, A. (2014). Opinion dynamics on interacting networks: media competition and social influence. *Nature, 4*, 1–7.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning II: Current research and theory* (2nd ed., pp. 64–99).
- Schelling, T. C. (1971). Dynamic models of segregation. *The Journal of Mathematical Sociology, 1*(2), 143–186.
- Schelling, T. C. (2006). *Micromotives and Macrobehaviour*. W.W. Norton & Company.
- Staudinger, M. R., & Büchel, C. (2013). How initial confirmatory experience potentiates the detrimental influence of bad advice. *NeuroImage, 76*, 125–133.
- Watts, D. J. (2002). A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences of the United States of America, 99*(9), 5766–5771.
- Watts, D. J., & Dodds, P. S. (2007). Influentials , Networks , and Public Opinion Formation. *Journal of Consumer Research, 34*(4), 441–458.
- Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of “small-world” networks. *Nature, 393*, 440–442.
- Zanette, D. H., & Gil, S. (2006). Opinion spreading and agent segregation on evolving networks. *Physica D, 224*, 156–165.