

Information Flows, the Accuracy of Opinions, and Crashes in a Dynamic Network

Phillip Monin

Office of Financial Research

phillip.monin@ofr.treasury.gov

Richard Bookstaber

University of California

rbookstaber@gmail.com

The Office of Financial Research (OFR) Staff Discussion Paper Series allows members of the OFR staff and their coauthors to disseminate preliminary research findings in a format intended to generate discussion and critical comments. Papers in the OFR Staff Discussion Paper Series are works in progress and subject to revision.

Views and opinions expressed are those of the authors and do not necessarily represent official OFR or U.S. Department of the Treasury positions or policy. Comments are welcome, as are suggestions for improvements, and should be directed to the authors.

Information Flows, the Accuracy of Opinions, and Crashes in a Dynamic Network*

Phillip Monin[†] and Richard Bookstaber[‡]

March 2, 2017

Abstract

Markets coordinate the flow of information in the economy, aggregating it through the price mechanism. We develop a dynamic model of information transmission and aggregation in financial and other social networks in which continued membership in the network is contingent on the accuracy of opinions. Agents have opinions about a state of the world and form links to others in a directed fashion probabilistically. Agents update their opinions by averaging those of their connections, weighted by how long their connections have been in the system. Agents survive or die based on how far their opinions are from the true state. In contrast to the results in the extant literature on DeGroot learning, we show through simulations that for some parameterizations the model cycles stochastically between periods of high connectivity, in which agents arrive at a consensus opinion close to the state, and periods of low connectivity in which agents' opinions are widely dispersed. We add varying degrees of homophily through a model parameter called tribal preference and find that crash frequency is decreasing in the degree of homophily. Our results suggest that the information aggregation function of markets can fail solely because of the dynamics of information flows, irrespective of shocks or news. (*JEL* codes: D83, D85, Z13)

Keywords: social networks; DeGroot learning; dynamic network formation; information transmission; nonlinear dynamical systems; crashes

*We gratefully acknowledge helpful comments from Matthew Jackson, Stathis Tompaidis, Peyton Young, Simpson Zhang, and seminar participants at the Conference on Information Transmission in Networks at Harvard University and at the Office of Financial Research (OFR). Remaining errors are the responsibility of the authors alone. Views and opinions expressed are those of the authors and do not necessarily reflect those of the OFR or the U.S. Department of the Treasury.

[†]phillip.monin@ofr.treasury.gov. Office of Financial Research, U.S. Department of the Treasury, 717 14th St. NW, Washington, DC 20005.

[‡]rbookstaber@gmail.com. University of California, Office of the Chief Investment Officer, 1111 Broadway, Oakland, CA 94607.

1 Introduction and background

Information is the life blood of the financial system. Indeed, efficient markets theory and information economics have long argued that information aggregation is the key role of markets in the economy. Prices in well-functioning, competitive markets reflect all available information and lead to efficient resource allocations. The economy's ex-post performance is then dependent on how accurate the information that led to this allocation was. If the information aggregation function of markets breaks down, then market failures can occur. In this paper we show that information aggregation can fail endogenously due to the dynamics of information flows among agents.

Individuals in society have opinions on a wide variety of topics. Some are economic or financial in nature, such as the value of a particular good or service, while others concern politics, social norms, entertainment, and so on. Decisions and actions are affected by opinions. Therefore, studying the dynamics of information flows and the processes by which information is aggregated is essential to understanding the macroscopic characteristics of both social and economic networks, regardless of the subject matter of agents' opinions or the precise mechanisms by which agents interact. Accordingly, social network models abstract from these details, often proposing simple mechanisms for how agents interact and how they use information gained to update their beliefs. While highly stylized, these models can provide important insights on features observed in real networks.

A central object of study in social networks is the evolution of opinions, particularly whether agents' opinions converge to a consensus and, if so, what that consensus is. In models with a fixed number of agents and a fixed network topology, a typical result is that convergence to a consensus is obtained in the long run, regardless of whether agents update rationally using Bayes's rule or naively by, for example, averaging the opinions of others (see, for example, Acemoglu and Ozdaglar (2011)). But what if *accuracy* matters, in that agents with inaccurate opinions fail and exit the system with a higher probability than those with relatively more accurate opinions? Such a situation can arise, for example, in a business context. If accurate information provides more value than inaccurate information, and thus leads to higher expected profits, then over time businesses with inaccurate information will fail due to competitive pressures.

We develop a dynamic model of information transmission and aggregation in social networks in which continued membership in the network is contingent on the accuracy of opinions. Agents in the model have an opinion about some fixed state of the world and form links to other agents probabilistically. Information is shared across connections and agents update their opinions according to a DeGroot learning rule,

that is, by continuously averaging those of their connections. The weights agents place on others' opinions depend on how reliable their information is judged to be. The reliability of an agent's opinion is proportional to how long it has been in the system. After updating their opinions, agents survive to the next period with a probability depending on how far their opinions are from the state of the world. Agents that do not survive fail and exit the system, and their information is lost. A failed agent's connections are severed, and the agent is replaced by an unconnected agent with random opinions.

Our model naturally incorporates DeGroot learning in a dynamic model of network formation in which the network topology dynamically depends on how accurate agents' opinions are with respect to an exogenous state of the world.

While our model is an abstract social networks model with no reference to specific applications, one can apply it in financial settings. One example concerns leverage (see section 3). The optimal level of leverage in the economy is unknown and, given that the economy is large and complex, it is reasonable to assume that financial institutions update their opinions on leverage by averaging those of other institutions, weighted by how long they have been in business. Moreover, there are consequences to having the wrong level of leverage: leverage levels too high or too low will lead to underperformance relative to peers, and possibly eventual failure due to competitive pressures. Another financial example concerns equity analysts, who publish forecasts on the performance of companies. Analysts who provide more reliable signals of a company's future performance develop reputations and are trusted more than other analysts. Analysts who do not produce reliable signals may leave the profession.

Our main result is that certain parameterizations of the model produce cyclical dynamics of information sharing. That is, the system alternates between a state of high connectivity, in which agents have achieved consensus opinions that are close to the state of the world, and a state of low connectivity, in which agents' opinions are widely dispersed. These cycles are stochastic in that they occur irregularly and last for varying amounts of time. This result is novel in the literature on DeGroot learning, directly contrasting with the asymptotic convergence of opinions to a consensus that is typical of other models (see, among others, DeGroot (1974), DeMarzo et al. (2003), and Golub and Jackson (2010)).

In the classical DeGroot (1974) model, agents with opinions about a state of the world are connected in a static network. Agents iteratively update their opinions by repeatedly taking weighted averages of the opinions of their connections. The main result from the literature is that, under mild assumptions on the nature of the network of connections, agents in the long run converge to a common opinion or consensus.

This is an intuitive result; if opinions are ordered, then agents with relatively high (low) opinions will have neighbors with low (high) opinions, and these will average out with the repeated updating over time.

DeMarzo et al. (2003) argue that DeGroot updating is a boundedly rational form of learning, as opposed to fully rational (Bayesian) forms of learning. Our motivation for using DeGroot updating rather than Bayesian learning is based on both conceptual and empirical considerations. Social networks are often large and extremely complex, which can render fully rational Bayesian approaches infeasible, as they may be too demanding on the agents (Golub and Jackson (2010) and Acemoglu and Ozdaglar (2011)). Boundedly rational or naive learning methods, such as DeGroot updating, might then be viewed as more realistic behavioral rules. Moreover, recent empirical evidence supports the view that naive learning methods such as DeGroot updating more realistically reflect the actual behaviors of agents than Bayesian learning methods (Mueller-Frank and Neri (2013), Grimm and Mengel (2014), and Chandrasekhar et al. (2015)).

In our model, a given agent forms a directed link with another agent in a given period according to an exogenous probability that is common across agents. For each agent at most one new link can be formed in a given period. Information is transmitted along all connections and, at any fixed time, each agent updates its opinion by taking a weighted average of those of its connections. The weight for each connection is proportional to its age, which is the number of consecutive periods for which it has been in the system. Links are destroyed when agents exit the system, i.e. when they “die” or “fail,” which happens with a probability depending on how far their opinions are from the state of the world. There is also a random component of agent failure in our model, irrespective of proximity to the true state (see also Jackson and Watts (2002) and Staudigl (2013)). This random component reflects the notion that some agents may simply leave the group regardless of their opinion for idiosyncratic reasons, such as death or moving on with their lives. Thus the so-called interaction matrix, which in DeGroot learning contains weights that agents place on the opinions of others in the averaging process, is time-varying, as agents continually reassess the relative trustworthiness of the information provided by others. This contrasts with the standard DeGroot model in which the interaction matrix is assumed constant over time. In addition, unlike Hegselmann and Krause (2002, 2005), Mirtabatabaei and Bullo (2012), and Weisbuch et al. (2002), who model time-varying interaction matrices with weights depending on similarity of opinions, agents in our model connect and transmit information in a directed fashion according to an exogenous probability and weight information depending on how long connections have been in the system.

Our paper also contributes to the literature on dynamic network formation (see, among others, Bala and Goyal (2000), Watts (2001), Jackson and Watts (2002), Jackson and Rogers (2007), Skyrms and Pemantle (2009), and Staudigl (2013)). Models in this literature propose and analyze mechanisms of dynamic network formation in which agents form and sever links through time. Bala and Goyal (2000) study directed communication networks in which the link formation process is governed by agents weighing costs of maintaining links against their benefits. Watts (2001) studies convergence properties in a network with a deterministic dynamic, while Jackson and Watts (2002) consider a wider class of models under stochastic evolution. Staudigl (2013) considers a general model of an interaction game in which the interaction probabilities are time-varying functions of the players' actions. While agents in these models form and sever links strategically, our agents have no direct control over how links are formed and severed. This is consistent with their naive DeGroot updating learning rule and has a similar motivation: in large and complex networks, it is reasonable to think that agents are limited in their ability to assess other agents and the ramifications of connecting to them prior to learning about them. Moreover, as there is no cost to connecting in our model, agents should want to connect to as many other agents as possible. It is only after the link has been made that the trustworthiness of an agent's information is assessed based on the proxy of how long the agent has been in the system.

After setting the model on a foundation for analysis by proving that it is mathematically represented as an ergodic Markov chain (see section 4), we show through simulations that the dynamics of the network can be the source of cycles of information sharing. That is, under a range of parameterizations, the model produces cycles of convergence and divergence of opinions, which occur despite holding the state of the world fixed. A cycle begins with the system in a state of high interconnectivity in which agents have achieved nearly consensus opinions that are close to the state of the world. Over time some agents die, either because their opinion is not precisely equal to the state of the world or because of exogenous reasons. The connections of dead agents are severed and they are replaced by naive ones with relatively inaccurate opinions. If such agents are able to connect to others, thereby spreading their relatively poor information, then their poor information can propagate quickly to the rest of the agents, causing the opinions of the rest of the agents to drift away from the state of the world. This can precipitate more deaths and the process can feed back onto itself, eventually causing a nearly complete collapse of the network. This is the process by which the state transitions to one of low connectivity in which agents' opinions are widely dispersed. We call such a transition from a state of high to low connectivity a

crash. Intuitively, the crash occurs because a very old interconnected agent with good information dies and is replaced by a relatively uninformed agent, who then spreads its poor information through the links it forms. Eventually a period of *rebuilding* occurs in which connections are reformed and the distribution of opinions narrows, arriving at a high-connectivity state. The cycle then repeats.

We continue our study by running a comparative statics analysis on the per-period probability of connection and the exogenous probability of failure. The results suggest that the level curves of crash frequency are approximately linear in these parameters.

Next, we extend the model to incorporate homophily, which is the widespread finding in the sociology and social networks literature that agents associate more with other agents that are similar to them and less with those that are not. Homophily has been observed across many agent characteristics, including age, gender, race, education level, profession, political affiliation (see McPherson et al. (2001) for a review). Since differential preferences for connecting to agents affects the link formation process, homophily influences the flow and aggregation of information. Golub and Jackson (2012) present and analyze a model with homophily and find that the time it takes to reach consensus in static models in which agents update beliefs using an averaging process such as DeGroot updating is increasing and convex in homophily.

We analyze the effect of homophily on network stability in a simple extension wherein we introduce the notions of tribe and tribal preference. A *tribe* is a pre-specified subset of the agents that are alike in some way that affects their propensity to connect with one another instead of the general populace. Such a binary grouping of agents as either “us” or “them” affecting the linking process has been shown in some networks to be fairly accurate (see, for example, Marsden (1987, 1988)). Conditional on connecting in a given period, a member of the tribe connects to another member of the tribe with a certain probability that we call *tribal preference*. If tribal preference is low, then agents connect randomly to other agents and do not discriminate among possible sources of information. Conversely, if tribal preference is high, then agents trust members of their tribe more and prefer to collect information from them.

We show that the dynamics of the extended model are governed by the relationships among tribal preference, the per-period probability of connection, and the exogenous probability of failure. We find that, not surprisingly, within-tribe connectivity increases as tribal preference increases. We also find that crash frequency is decreasing in tribal preference. For low levels of tribal preference, where all agents in the tribe are connecting indiscriminately among the larger population, the agents are exposed to a greater number of agents. The more agents there are, the higher the probability that at least one agent dies (either for exogenous reasons or because of imperfect opinions). If

a dead agent is replaced by a new agent whose information is poor relative to other agents but also accurate enough in an absolute sense to survive and connect, then in a highly connected state this relatively poor information can propagate widely, reducing the accuracy of others' information, leading to more deaths and eventually a crash. Conversely, for high levels of tribal preference, all agents within the tribe are mainly connecting to themselves. As the tribe is a proper subset of the general population, this leads to fewer chances for a dead agent's replacement to infect the tribe and cause a crash. Over periods of the same length, lower tribal preference thus leads to a greater number of crashes.

2 A dynamic model of information flows

Agents in our model operate in a discrete-time environment over an infinite horizon. An agent's objective in a given period is to maximize the probability of surviving to the next period. Each agent has an opinion $k \in (0, 1) \cap \mathbb{Q}$ about a state of the world $k^* \in (0, 1) \cap \mathbb{Q}$, where \mathbb{Q} is the set of rational numbers. The true state k^* is fixed and unknown to all of the agents. In a financial context, for example, k could be the (reciprocal of the) agent's leverage ratio, in which case k^* would be the (reciprocal of the) optimal leverage ratio in the economy. An agent's probability of survival to the next period depends on how close its opinion k is to k^* . To improve their chances of survival, agents connect and share information that they then use to update their opinions.

Specifically, the model contains $N \geq 2$ agents. Time is discrete and is denoted by $t \in \mathbb{N} \cup \{0\}$, where \mathbb{N} is the set of natural numbers. At initial time $t = 0$, opinions are drawn independently across agents from a uniform distribution on $(0, 1) \cap \mathbb{Q}$. Let $k_{i,t} \in (0, 1) \cap \mathbb{Q}$ be the opinion of agent i at time t . In addition to its opinion $k_{i,t}$, each agent is aware of its age $a_{i,t} \in \mathbb{N}$, which is the number of consecutive periods for which it has been alive. Denote by $\mathbf{k}(t)$ and $\mathbf{a}(t)$ the N -vectors of opinions and ages of agents, respectively, at time t .

Agents adjust their opinions to improve their chances of survival to the next period. This adjustment is based on information learned by connecting to other agents. At birth, agents have zero connections. Then, in each period in which it is alive, with probability $p^c \in (0, 1)$, a given agent i creates a directed link with another, randomly chosen agent j . Information transmission along a connection between agents is one-way, that is, information flows from agent j to agent i , with agent i learning agent j 's k -value and age, $(k_{j,t}, a_{j,t})$. An agent's information set at any point in time therefore contains its own opinion and age as well as the opinions and ages of all of its in-links (but not

of its out-links). Using this information, agents update their opinions according to a DeGroot-type learning rule,

$$\mathbf{k}(t+1) = \mathbf{W}(t)\mathbf{k}(t), \quad (1)$$

where $\mathbf{W}(t) = (w_{ij}(t))$ is the weight matrix with entries defined by

$$w_{ij}(t) = \begin{cases} a_{j,t} / \sum_{k \in N(i,t)} a_{k,t}, & j \in N(i,t) \\ 0, & j \notin N(i,t) \end{cases} \quad (2)$$

and $N(i,t)$ is the union of the in-links of agent i and itself at time t . Each agents' updated opinion is thus a weighted average of the opinions of its in-links (and itself), weighted by age.¹

Survival depends on the distance between opinions and the state of the world $k^* \in (0,1) \cap \mathbb{Q}$. The closer an agent's opinion is to k^* in a given period, the higher the probability that the agent will live to the next period. Formally, if $k_{i,t}$ is the opinion of agent i at time t , then the probability $p_{i,t}^d$ that the agent will die at time t is given by

$$p_{i,t}^d = \beta_0 + (k_{i,t} - k^*)^2, \quad (3)$$

where β_0 is a positive parameter of the model chosen such that $p_{i,t}^d \in (0,1)$. The positivity of β_0 implies that there is a positive probability of death independent of the accuracy of the agent's opinion. That is, there are exogenous shocks where agents die with a probability independent of the relationship between their opinions and the state of the world.

When an agent dies, its connections with all other agents are severed immediately, and thus the information embedded in its opinion and age are lost to all other agents. A dead agent is replaced in the next period by a new agent with no initial connections and initial opinion uniformly drawn from $(0,1) \cap \mathbb{Q}$. Thus the population of agents remains fixed.

2.1 Optimization and averaging based on age

The objective of each agent is to survive to the next period. In that sense, the model could be viewed as a series of one-period models. Here we investigate whether

¹Agents do not take into account the paths or evolution of opinions of their present or past connections in their formation of opinions. This information is embedded in previous opinions. Agents are thus memoryless. There are various ways to incorporate memory. For example, one could exponentially weight past information. But this introduces additional parameters and complexity. Moreover, the "perfect memory" case, where agents perfectly remember their entire past, has qualitatively similar results to those of the present model.

an agent's opinion, which it updates heuristically based on a weighted average of the opinions of its connections with weights proportional to its connections' ages, is optimal with respect to its objective. We show that, under some assumptions, an agent's opinion in any given period is approximately its maximum likelihood estimate of the state of the world k^* . Weighting based on age is thus approximately optimal.

Proposition 2.1. *Fix an arbitrary time $t \in \mathbb{N}$ and an arbitrary agent $i \in \{1, \dots, N\}$. Let $\{(k_{j,t}, a_{j,t})\}_{j=1}^n$ be the information set for agent i at time t assembled from its in-links (and itself) at t . Consistent with the model, assume that agent i is unaware of the history of all agents' opinions and actions prior to t , and that, consequently, it assumes all of its connections' opinions prior to t are equal to their values at time t . Then,*

$$k_{i,t+1} = \frac{\sum_{j=1}^n a_{j,t} k_{j,t}}{\sum_{j=1}^n a_{j,t}} \quad (4)$$

is the first-order approximation of the agent's maximum likelihood estimate of k^* .

Proof. To simplify notation we omit the t -subscripts. Using (3), the assumptions imply that the probability that an arbitrary agent $j \in \{1, \dots, N\}$ has age a_j is given by

$$\mathbb{P}(\text{agent } j \text{ has age } a_j) = (1 - \beta_0 - (k_j - k^*)^2)^{a_j}.$$

Moreover, the joint probability that all agents have their respective ages is given by the product of the above over all agents, which results in the likelihood function

$$f(k_j, a_j; j = 1, \dots, n \mid k^*) = \prod_{j=1}^n (1 - \beta_0 - (k_j - k^*)^2)^{a_j}.$$

Let \mathcal{L} be the log-likelihood function, i.e.

$$\mathcal{L} = \sum_{i=j}^n a_j \cdot \log(1 - \beta_0 - (k_j - k^*)^2).$$

While \mathcal{L} is continuously differentiable, maximizing it with respect to k^* requires finding the roots of a $(2n - 1)$ -degree polynomial. Instead, we recall the Taylor series

$$\log(1 - x) = - \sum_{n=1}^{\infty} \frac{x^n}{n} = -x - \frac{x^2}{2} - \frac{x^3}{3} - \dots,$$

and use it to find the first-order Taylor expansion of \mathcal{L} , namely

$$\mathcal{L}' = - \sum_{j=1}^n a_j (\beta_0 + (k_j - k^*)^2).$$

Then, differentiating \mathcal{L}' with respect to k^* , we obtain

$$\frac{\partial \mathcal{L}'}{\partial k^*} = 2 \sum_{j=1}^n a_j (k_j - k^*).$$

Solving $\frac{\partial \mathcal{L}'}{\partial k^*} = 0$ yields the critical point given by (4) above, which is easily seen to be the unique maximum of \mathcal{L}' . \square

3 Financial interpretations

While our model of information flows among agents is abstract, there are several natural applications of it in a financial context. One application concerns leverage. In its simplest form, leverage refers to the use of borrowed funds to obtain more economic exposure than would be possible using only one's capital. One motivation is to enhance returns. Suppose you borrow \$80 against capital of \$20 and invest the sum in some asset. You own an asset worth \$100 using capital of \$20 and are thus leveraged 5-to-1, with a leverage ratio of 5. If you sell the asset when it appreciates 5% to \$105 and then pay back your loan of \$80 (assuming zero interest for simplicity), then you are left with \$25, a 25% return on your capital of \$20. Thus a 5% return on the asset implies a 25% return on your capital if your leverage ratio is 5. The risk is that leverage magnifies negative returns as well. If you sell the asset after it depreciates 5% to \$95 and pay back your loan, you have \$15, a -25% return on your capital.

Geanakoplos (1997, 2003, 2010) argues that leverage is an equilibrium state variable in the economy. However, the optimal level of leverage, that is, the equilibrium level of leverage that is a consequence of agents' portfolio choice problems, is unknown. Because of its returns-magnifying property, too much leverage can risk solvency. But in a competitive economy, too little leverage can lead to underperformance.

Our model can be applied in this setting. Consider a set of financial institutions, such as banks or hedge funds. The financial institutions are the agents in the network. Each financial institution's opinion, k , is the reciprocal of its leverage ratio, and the state of the world, k^* , is the unknown, reciprocal of optimal leverage ratio.² Connections are formed between institutions in the course of doing business. For example, a direct connection from fund j to bank i is formed as bank i learns fund j 's leverage when the fund applies for financing from the bank. Alternatively, connections between institutions arise as a result of negotiations about collateral agreements, where leverage might be used to assess credit risk of counterparties. The limit on connections to

²An institution's leverage ratio is generally in $[1, \infty)$, so that the reciprocal of its leverage ratio is in $(0, 1]$.

one per period with a probability of connecting p^c reflects technological capacity for handling new business and information processing.

Most economic and financial models posit that agents act rationally. But in a large network like the global financial system, where fully rational learning may be infeasible or impractical, the DeGroot updating rule, in which the institution updates its leverage based on the choices of other institutions weighted by how long they have survived, is a reasonable behavioral rule.

The link destruction process is also naturally interpreted in the context of leverage. Agent death represents going out of business because of a suboptimal amount of leverage, resulting from poor performance relative to peers and ensuing competitive pressures. The parameter β_0 reflects going out of business for reasons unrelated to leverage, such as operational failures. Finally, birth of agents upon death represents new entrants.³

Another application of our model concerns equity analysts. Equity analysts publish forecasts of earnings and other performance-related measures of companies. In the application of our model, the agents in the network are these equity analysts, who for simplicity we will assume are all covering the same company. Each analyst's opinion k represents its opinion on the future performance of the company, and k^* represents the unknown, true future performance. As there are typically many analysts across the industry that analyze a given company, random connections among analysts are made when their opinions are published or when they otherwise communicate. When connections are made, it is reasonable to assume that analysts update their opinions, using age or length of time in the industry as a proxy for another analyst's reputation and thus trustworthiness. Analysts exit the network and sever their ties for exogenous reasons such as retirement or moving on to a different job, as well as possibly for inaccurate opinions.

Finally, another application of our model concerns the network of information flows that results from the piecing together of disparate information in what is called the mosaic approach to investing. The agents of the network are investors, who evaluate the prospects of a company by piecing together things such as the fragility of supply chains, the availability of funding to meet capital expenditures, new entrants into the field, disruption from technological innovations, and changes in consumer preferences. This information will come from different agents, and the accuracy of the information

³The model assumes agents are homogeneous in their behavioral rule, objective, probability of connection, prior distribution of opinions, exogenous probability of death, and so on. These assumptions are idealizations made to simplify exposition, implementation, and analysis, as well as possible calibration of the model to data. In reality agents are heterogeneous, in these respects and others.

will dictate the probability of survival for the enterprise. The state of that enterprise will in turn be one part of the mosaic that is being filled out by other agents.

4 View of the model as a Markov chain

We show that the model can be viewed in a mathematical context as an ergodic Markov chain. Such a representation enhances our understanding of the model and provides a foundation for our subsequent simulations.

4.1 Relevant background on Markov chains

Before elaborating on the connection between the model and ergodic Markov chains, we need to introduce some terminology. We will only introduce what we need. The analysis of Markov chains is well-developed; comprehensive treatments on the subject and its applications can be found in Kallenberg (2002), Pardoux (2008), and Ibe (2009), among others.

A *Markov chain* is a stochastic process $\mathcal{X} = \{X_t; t \in \mathbb{N}\}$ such that, for all $t \in \mathbb{N}$, the conditional distribution of X_{t+1} given the history X_0, X_1, \dots, X_t equals its conditional distribution given the present, i.e. given X_t . The *state space* S associated to a Markov chain is the range of the stochastic process \mathcal{X} , and a Markov chain is called *finite (countable)* if its state space is finite (countable). To each Markov chain one associates a *transition probability* function $p: \mathbb{N} \times S \times S \rightarrow [0, 1]$ for which $p(t, i, j)$ is the probability of moving from state i to state j at time t . The chain is called *time-homogeneous* if this probability is independent of the time t , i.e. if the function p is independent of t .

The analysis of Markov chains begins with the classification of states along a number of dimensions. Suppose a Markov chain has initial state $X_0 = x$ and define $T_x = \inf\{t \in \mathbb{N}: X_t = x\}$, that is, the first time the chain returns to x . If the probability (conditional on the chain starting at x) that T_x is finite is one, then the state x is said to be *recurrent*. Otherwise, x is called *transient*. For a recurrent state x , if the stronger statement that the expected value of T_x (conditional on the chain starting at x) is finite holds, then x is called *positive recurrent*. Let R_x be the set of integers n for which the probability, conditional on starting at x , of returning to x in n periods is positive. The state x is called *aperiodic* if the greatest common divisor of R_x is one.

When all states in the state space satisfy a certain property, we say that the chain itself satisfies the property. For example, if all states are aperiodic then we say the chain is itself aperiodic. It turns out this type of inheritance is common with Markov chains. Recurrence, for example, is a class property of the Markov chain, that is, the state space

can be partitioned into a set of equivalence classes wherein each class has the property that all states within it are either recurrent or transient. The equivalence relation \sim producing this partition is defined as follows: two states *communicate*, written $x \sim y$, if they are accessible from each other, where y is *accessible* from x , and denoted by $x \rightarrow y$, if there is a (strictly) positive probability of going from state x to state y in finite time. A Markov chain is said to be *irreducible* if its state space under the equivalence relation \sim can be partitioned into a single equivalence class. An irreducible, positive recurrent, and aperiodic Markov chain is said to be *ergodic*.

An ergodic Markov chain has special asymptotic properties. For every $t \in \mathbb{N}$ let $f_t(x|x_0)$ denote the *empirical frequency* that describes how often the chain has been in state x through time t , given that it was initially in state x_0 . An ergodic Markov chain has the property that this empirical frequency distribution converges almost surely to a unique limiting distribution independent of the initial state.⁴ That is, there exists a probability distribution λ on S such that

$$\lim_{t \rightarrow \infty} f_t(x|x_0) = \lambda(x),$$

where the convergence is almost surely. In fact, a stronger statement can be made. For arbitrary states $x, y \in S$, it holds with probability one that

$$\lim_{t \rightarrow \infty} \mathbb{P}(X_t = x | X_0 = x_0) = \lambda(x).$$

Thus for an ergodic chain, if t is sufficiently large, then the probability of being in a state x at time t is approximately equal to the probability $f_t(x|x_0)$ of being in state x through time t .

4.2 Reformulation of the model

We reformulate the model described in section 2 as a Markov chain \mathcal{X} on a countable state space whose transition probability is characterized by an algorithm that we specify below. The Markov chain \mathcal{X} is governed by the following set of parameters.

Assumption 1. We have the following assumptions on the parameters of \mathcal{X} :

- (a) Number of agents: $N \in \mathbb{N}$
- (b) State of the world: $k^* \in [\frac{1}{2}, \sqrt{1 - \beta_0})$
- (c) Probability of connecting: $p^c \in (0, 1)$
- (d) Death probability intercept: $\beta_0 \in (0, 1)$

⁴Some define ergodicity of a Markov chain as precisely this, i.e. that the limiting distribution of the empirical frequency is independent of the initial state.

The state of the Markov chain \mathcal{X} at any given time is the set of opinions and ages of all the agents, together with the adjacency matrix describing the connections among agents.

Assumption 2. We have the following assumptions on the state variables of \mathcal{X} :

- (a) k -value vector: $\vec{k} = (k_1, \dots, k_N) \in (\mathbb{Q} \cap (0, 1))^N$
- (b) Age vector: $\vec{a} = (a_1, \dots, a_N) \in \mathbb{N}^N$
- (c) Adjacency matrix: $\vec{M} = (m_{ij})$, where $m_{ij} = 1$ if there is a directed link from i to j , i.e. i is an in-link of j . We adopt the convention that $m_{ii} = 1$ for all i . We write $\vec{M} = [\vec{m}_1 \vec{m}_2 \cdots \vec{m}_N]$, where each \vec{m}_i is the column of \vec{M} . The set of in-links of agent i correspond to the nonzero entries of the column vector \vec{m}_i .

The following algorithm is a mathematical description of the model described in section 2. The algorithm describes the transition probability function from one state to another in the Markov chain \mathcal{X} .

Algorithm: Given $(\vec{k}_{t-1}, \vec{a}_{t-1}, \vec{M}_{t-1})$ at time t .

A1. **Check Death:** Simultaneously, determine if each agent $i, i = 1, \dots, N$, dies or survives. The probability that agent i dies is given by

$$\mathbb{P}(\text{agent } i \text{ dies at } t) = \beta_0 + (k_{i,t-1} - k^*)^2,$$

where $k_{i,t-1}$ is the opinion of agent i .

(a) If agent i dies, then set

- $\hat{k}_{i,t-1} = U$, where U is uniformly drawn on $(0, 1) \cap \mathbb{Q}$;
- $a_{i,t} = 1$; and
- $m_{ii} = 1$ and $m_{ij} = m_{ji} = 0$ for all $j \neq i$.

(b) If agent i survives, then increment its age:

- $a_{i,t} = a_{i,t-1} + 1$.

A2. **Connect:** Agents connect to other agents. Let \bar{N}_i^c denote the set of agents to whom agent i is not connected, i.e.

$$\bar{N}_i^c = \{A_j : m_{ji} = 0\}.$$

If $\bar{N}_i^c = \emptyset$, then there are no agents left to whom agent i can connect. If \bar{N}_i^c is nonempty, then agent i connects to a member of it with probability p^c . If agent i connects to agent j , then set $m_{ji} = 1$.

A3. **Update:** Agents updates their opinions.

$$k_{i,t} = \sum_{A_j \in \bar{N}_i} w_{j,t} k_{j,t-1}, \quad w_{j,t} = a_{j,t} / \sum_{A_k \in \bar{N}_i} a_{k,t}$$

Steps A1-A3 determine $(\vec{k}_t, \vec{a}_t, \vec{M}_t)$.

4.3 Discussion of attainable states

The above algorithm naturally determines the set of states that are attainable in the model. For example, each agent can connect to at most one other agent in any given period, and an attainable state $(\vec{k}, \vec{a}, \vec{M})$ must have the property that

$$a_i \geq \sum_{j=1}^N m_{ji}, \quad i = 1, \dots, N. \quad (5)$$

Moreover, it is easy to see that not every conceivable state satisfying (5) is attainable. For example, consider a 2-agent system and a state $(\vec{k}_t, \vec{a}_t, \vec{M}_t)$ at time t given by

$$\vec{k}_t = (1/4, 3/4), \quad \vec{a}_t = (1, 1), \quad \text{and} \quad \vec{M}_t = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}. \quad (6)$$

This state requires that both agents die at time t , that the first agent connects to the second agent following their rebirth, and that opinions update to $\vec{k} = (1/4, 3/4)$. Upon the deaths of the agents, the newly drawn opinions $(\hat{k}_{1,t-1}, \hat{k}_{2,t-1})$ of the two agents must satisfy

$$\begin{pmatrix} \frac{1}{4} \\ \frac{3}{4} \end{pmatrix} = \begin{pmatrix} \frac{1}{2} & 0 \\ \frac{1}{2} & 1 \end{pmatrix}' \begin{pmatrix} \hat{k}_{1,t-1} \\ \hat{k}_{2,t-1} \end{pmatrix},$$

which has the unique solution $\hat{k}_{1,t-1} = -\frac{1}{4}$ and $\hat{k}_{2,t-1} = \frac{3}{4}$. Since $\hat{k}_{1,t-1} \notin (0, 1) \cap \mathbb{Q}$, this state is not attainable.

Note that, as in the algorithm, we use the symbol $\hat{\cdot}$ to signify that the associated value of k represents a newly drawn opinion for a reborn agent. Thus $\hat{k}_{1,t-1}$ is the new draw for agent 1 after its death at time t (step A1 in the algorithm, before it updates its opinion in step A3), while $k_{1,t-1}$ is the opinion of agent 1 at the end of the previous period $t - 1$ (upon completion of step A3 at time $t - 1$, just prior to its death in step A1 at time t).

For more insight on attainable states, consider a 2-agent system and a state at time t with opinions vector $\vec{k}_t = (k_{1,t}, k_{2,t})$, age vector $\vec{a}_t = (3, 1)$ and adjacency matrix \vec{M}_t equal to the one in (6). This state requires that the first agent dies three periods prior (step A1 at time $t - 2$), the second agent dies in the most recent period, and the first agent then connects to the reborn second agent.

The opinions of the agents at time $t - 1$ must satisfy

$$\begin{pmatrix} k_{1,t} \\ k_{2,t} \end{pmatrix} = \begin{pmatrix} \frac{3}{4} & 0 \\ \frac{1}{4} & 1 \end{pmatrix}' \begin{pmatrix} k_{1,t-1} \\ \hat{k}_{2,t-1} \end{pmatrix}, \quad (7)$$

where again $\hat{k}_{2,t-1}$ denotes the newly drawn opinion for agent 2. This system has a unique solution given k_t since the matrix in (7) is invertible. If the solutions are not between 0 and 1, then the state is not attainable. If they are, then, in particular, the newly drawn opinion $\hat{k}_{2,t-1}$ for the second agent must equal $k_{2,t}$.

Observe that the second agent's opinions and connections prior to the period in which it dies are irrelevant. This is because when it dies its connections are lost, its age is reset, and its opinions are redrawn. Moreover, the second agent's opinions and connections prior to its death are also irrelevant to the first agent. To obtain an age of 3

at period t , the first agent must die at the beginning of period $t - 2$, at which time it is reborn with age 1 and newly drawn opinions. Agent 2 either lives or dies in period $t - 2$ and agents 1 and 2 either connect or do not. In period $t - 1$, agent 1 lives but agent 2 either lives or dies, and again, either the agents connect if they were previously unconnected, or they do not. None of this matters for the attainability of the state. For the sake of argument, suppose that agent 2 has an age of 5 at time $t - 2$ and that agent 1 connects to agent 2 in that period. Then

$$\vec{a}_{t-2} = (1, 5), \quad \vec{a}_{t-1} = (2, 6), \quad \text{and} \quad \vec{M}_{t-2} = \vec{M}_{t-1} = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}.$$

Since equation (7) shows the evolution from $t - 1$ to t , we must consider the evolution from $t - 3$ to $t - 1$. After some computations we obtain that

$$\begin{pmatrix} k_{1,t-1} \\ k_{2,t-1} \end{pmatrix} = \begin{pmatrix} \frac{1}{24} & 0 \\ \frac{23}{24} & 1 \end{pmatrix}' \begin{pmatrix} \hat{k}_{1,t-3} \\ k_{2,t-3} \end{pmatrix}, \quad (8)$$

which, in turn, has a unique solution that, if between 0 and 1, determines the newly drawn opinions, $\hat{k}_{1,t-3}$, of agent 1. However, this complicated setup is unnecessary when merely considering whether a state is attainable. The reason is the Markovian nature of the model. Another path that is equivalent from the perspective of attainability, i.e. that obtains opinions at time t equal to the right side of (7), is for agent 1 to be unconnected to agent 2 in the periods before the death of agent 2, thus with constant opinions. In this case, its newly drawn opinions, $\hat{k}_{1,t-3}$, for agent 1 after its death are equal to $k_{1,t-1}$ from (7).

If existence of agent 2 is irrelevant to agent 1 before the death of agent 1, then it is reasonable to think that the evolution of the model before the death of agent 1 is also irrelevant. This is true. The behavior of the model preceding three periods ago has no influence on the attainability of a state. This is because to obtain the age vector $\vec{a}_t = (3, 1)$ all agents *must* die at some point in the last three periods, which severs their connections. They must then be reborn with newly drawn opinions. Therefore, the opinions, ages, and connections in all preceding periods are irrelevant. With respect to attainability, agents are irrelevant before their deaths.

Generally, the set of possible paths to a state $(\vec{k}_t, \vec{a}_t, \vec{M}_t)$ is mainly constrained by the age vector \vec{a}_t and adjacency matrix \vec{M}_t . These, in turn, constrain the set of possible opinion vectors \vec{k}_t . To see this, consider that the components of the age vector \vec{a}_t set the times prior to the current period at which the agents must have died. If, for example, $\vec{a}_t = (5, 3, 2)$, then the first, second, and third agents must have died 5, 3, and 2 periods ago, respectively. By the above discussion, all periods before the period at which the

oldest agent (according to \vec{a}_t) died are irrelevant to the attainability of the state. Given this and (5), there are finitely many sequences of adjacency matrices that result in \vec{M}_t . Enumerate the possible paths of $\{\vec{a}_s, \vec{M}_s\}_{s=t+1-a(N)}^t$, where t denotes the current period and $a(N)$ denotes the maximal value of \vec{a} . Next, fix such a path and examine whether there exists a viable sequence $\{\vec{k}_s\}_{s=t+1-a(N)}^t$ of opinions. If so, then the state $(\vec{k}_t, \vec{a}_t, \vec{M}_t)$ is attainable by traversing this path, provided it has positive probability (the state could also be attainable along other paths). If no viable sequence of opinions can be found for all paths, then the state is not attainable. Observe that the 2×2 -matrices in (6), (7), and (8) depend on the age vectors and adjacency matrices in the intervening periods, but not on the opinions vectors. It is straightforward to show that, for periods in which deaths do not occur,

$$\begin{aligned}\vec{k}_t &= f(\vec{a}_t, \vec{M}_t) \vec{k}_{t-1} \\ &= f(\vec{a}_t, \vec{M}_t) f(\vec{a}_{t-1}, \vec{M}_{t-1}) \vec{k}_{t-2} \\ &= \dots \\ &= f(\vec{a}_t, \vec{M}_t) f(\vec{a}_{t-1}, \vec{M}_{t-1}) \dots f(\vec{a}_{t-s+1}, \vec{M}_{t-s+1}) \vec{k}_{t-s} \\ &= F\left(\{\vec{a}_n\}_{n=t-s+1}^t, \{\vec{M}_n\}_{n=t-s+1}^t\right) \vec{k}_{t-s},\end{aligned}$$

where f is the weight matrix \mathbf{W} in (2) and F is the composition of weight matrices. However, such a recursion does not hold in periods in which agent deaths occur. Instead, one appeals to backward induction. Let $a_{(1)}$ denote the minimal age in \vec{a}_t . The last agents die in period $t+1-a_{(1)}$, at which time they are reborn with newly drawn opinions. Opinions at time $t+1-a_{(1)}$ must then satisfy

$$\vec{k}_t = F\left(\{\vec{a}_n\}_{n=t+2-a_{(1)}}^t, \{\vec{M}_n\}_{n=t+2-a_{(1)}}^t\right) \hat{\vec{k}}_{t+1-a_{(1)}},$$

where $\hat{\vec{k}}_{t+1-a_{(1)}}$ indicates the vector including these newly drawn opinions and with other components given by those in $\vec{k}_{1,t+1-a_{(1)}}$. The matrix F , which depends on the fixed sequence of age vectors and adjacency matrices under consideration, will determine whether there is a solution such that all opinions are between 0 and 1. If not, then the state is not attainable along this path. We then discard this path, consider a different path of $\{\vec{a}_s, \vec{M}_s\}_{s=t+1-a(N)}^t$, and start again. If, on the other hand, a viable solution is found, then we continue. In this case, let $a_{(2)}$ denote the second-least value in \vec{a}_t . The corresponding agents die in period $t+1-a_{(2)}$, at which time they are reborn with newly drawn opinions. Opinions at time $t+1-a_{(2)}$ must satisfy

$$\vec{k}_{t+1-a_{(1)}} = F\left(\{\vec{a}_n\}_{n=t+2-a_{(2)}}^{t+1-a_{(1)}}, \{\vec{M}_n\}_{n=t+2-a_{(2)}}^{t+1-a_{(1)}}\right) \hat{\vec{k}}_{t+1-a_{(2)}}.$$

Again, if a solution exists and is such that all opinions are between 0 and 1, then continue. If not, then the state is not attainable along the specific path. Continue in like manner until all agents are considered. If viable solutions exist for all times at which deaths occur, then the state is attainable along the path.

4.4 Ergodicity

Finally, we prove that the Markov chain \mathcal{X} is ergodic. Ergodicity of the Markov chain representation of our model provides a foundation for the use of simulations to accurately estimate population parameters and other quantities of interest. In essence, a single time series of sufficient length is representative of the entire underlying process.

Theorem 4.1. *The Markov chain \mathcal{X} is irreducible, aperiodic, and positive recurrent. It is thus ergodic.*

Proof. First, we show that \mathcal{X} is irreducible. That is, if $(\vec{k}, \vec{a}, \vec{M})$ and $(\vec{k}', \vec{a}', \vec{M}')$ are attainable (having dropped the t -subscripts), then

$$(\vec{k}, \vec{a}, \vec{M}) \rightarrow (\vec{k}', \vec{a}', \vec{M}').$$

It suffices to exhibit a path starting at state $(\vec{k}, \vec{a}, \vec{M})$ that has a positive probability of reaching the target state $(\vec{k}', \vec{a}', \vec{M}')$ in a finite number of periods. Write the target age vector $\vec{a}' = (a'_1, \dots, a'_N)$ and let t^* be the maximal age in the target state, that is, $t^* = \max_i a'_i$. By the discussion on attainability (see section 4.3), there is a path $\{\vec{a}_t, \vec{M}_t\}_{s=t+1-t^*}^t$ and an associated viable sequence $\{\vec{k}_s\}_{s=t+1-t^*}^t$ of opinions such that the state is attainable. This sequence of states has positive probability since $\beta_0 > 0$ and $p^c \in (0, 1)$. (Here and throughout, we consider a transition to a state $(\vec{k}, \vec{a}, \vec{M})$ to have positive probability if the transition to $(\vec{k} \pm \epsilon \mathbf{1}, \vec{a}, \vec{M})$, where $\mathbf{1}$ is the N -vector of ones, has positive probability for any $\epsilon > 0$.)

Next, we show that \mathcal{X} is aperiodic. Since \mathcal{X} is irreducible, it suffices to show that there exists one state that is aperiodic (see, for example, Pardoux (2008), Lemma 6.2). That is, it suffices to find one state s for which $\mathbb{P}(X_{t+1} = s | X_t = s) > 0$ for an arbitrary $t \in \mathbb{N}$. Consider the state $s_0 = (\vec{k}, \vec{a}, \vec{M})$ where $\vec{a} = (1, 1, \dots, 1)^\top$, \vec{M} is the N -dimensional identity matrix, and $\vec{k} = k \in (0, 1)^N$ is fixed but arbitrary. The event that $X_{t+1} = s_0$ given that $X_t = s_0$ corresponds to the event in which every agent dies, no connections are made, and the vector of opinions remains the same. This event has positive probability since agent deaths are independent, the probability of death of a given agent in a given period is strictly between zero and one, the probability of an agent connecting in a given period is less than one, and the event that k -values remain the same has positive probability.

Finally, we show that \mathcal{X} is positive recurrent. Since \mathcal{X} is irreducible, it suffices to show that there exists one state that is positive recurrent. We will show s_0 is positive recurrent, i.e. that

$$m(s_0) = \sum_{t=1}^{\infty} t \cdot \mathbb{P}(T_{s_0} = t) < \infty,$$

where

$$T_{s_0} = \inf\{s \geq 0: X_s = s_0 \mid X_0 = s_0\}.$$

Recall that agent deaths are independent across agents and time (agent interactions affect k -values, but, for example, the conditional probability that agent i dies conditional on the death of agent j is just the unconditional probability that agent i dies). Using Assumption 1 and (3), it is easy to show that the probability of death for a given agent i at a given time t satisfies

$$0 < \beta_0 \leq p_{i,t}^d \leq \beta_0 + (k^*)^2 < 1.$$

Therefore, the probability that at least one agent survives in a given period satisfies

$$\begin{aligned} \mathbb{P}(\text{at least one agent survives}) &= 1 - \mathbb{P}(\text{all agents die}) \\ &\leq 1 - \beta_0^N. \end{aligned}$$

Next, observe that the event that $T_{s_0} = t$ is a subset of the event in which at least one agent survives in each of the first $t - 1$ periods and *all* agents die in period t . By the above, we have that

$$\begin{aligned} \mathbb{P}(T_{s_0} = t) &\leq \mathbb{P}(\text{at least one agent survives in first } t - 1 \text{ periods}) \\ &\quad \times \mathbb{P}(\text{all agents die at time } t) \\ &\leq (1 - \beta_0^N)^{t-1} (\beta_0 + (k^*)^2) \\ &\leq c^t, \end{aligned}$$

where $c := \max\{1 - \beta_0^N, \beta_0 + (k^*)^2\} < 1$. Therefore,

$$m(s_0) = \sum_{t=1}^{\infty} t \cdot \mathbb{P}(T_{s_0} = t) \leq \sum_{t=1}^{\infty} tc^t,$$

which converges by the ratio test. □

5 Simulations: Cycles of consensus and dispersed opinions

The analysis of models of opinion dynamics usually appeals to results in the overlapping fields of graph theory, Markov processes, and dynamical systems. The object of study is the dynamic properties of the opinions updating process (1), particularly asymptotic convergence results. Static models of DeGroot updating, where the opinions updating process \mathbf{W} from equation (1) is constant, are linear dynamical systems and have been treated at length. Golub and Jackson (2010), for example, produce a general result establishing a set of mild assumptions that lead to convergence of opinions in the static DeGroot model.

Analytic or asymptotic results are more difficult to obtain in models such as ours, where \mathbf{W} is time-varying. This is especially true if the opinions of the agents affect the network structure and, in turn, the network structure affects the opinions of agents. The feedback between the two creates a nonlinear dynamic. Hegselmann and Krause (2002, 2005) (see also Lorenz (2005, 2006a,b)) analyze models with time-varying \mathbf{W} in which agents dynamically form links to other agents with similar opinions. These authors obtain some convergence results but unfortunately their methods do not apply to our model.

An alternative approach is to analyze the model using Monte Carlo simulations. Because of the ergodic nature of the model when viewed as a Markov chain (see section 4), Monte Carlo simulations lead to results that are asymptotically valid. Hence we can be assured that the behavior of the system we observe in the simulation will converge to its true behavior and will not depend on the particular initial condition used to begin the simulation.

Table 1: Baseline simulation parameters

Parameter	Value
N	25
k^*	0.6
p^c	5×10^{-3}
β_0	2×10^{-5}

From the model description in section 2 we deduce that our model is parameterized by the 4-tuple (N, k^*, p^c, β_0) , where $N \in \mathbb{N}$ is the number of agents, $k^* \in [\frac{1}{2}, \sqrt{1 - \beta_0})$ is the state of the world, $p^c \in (0, 1)$ is the common probability of an agent connecting to another agent in a given period, and $\beta_0 \in (0, 1)$ is the exogenous probability of

death (see (3)). Table 1 lists the baseline parameter values used for the simulations, which were arrived at by experimentation. A sensitivity analysis of the parameters suggests that the qualitative results that we describe are robust to modifications (see also section 6).⁵

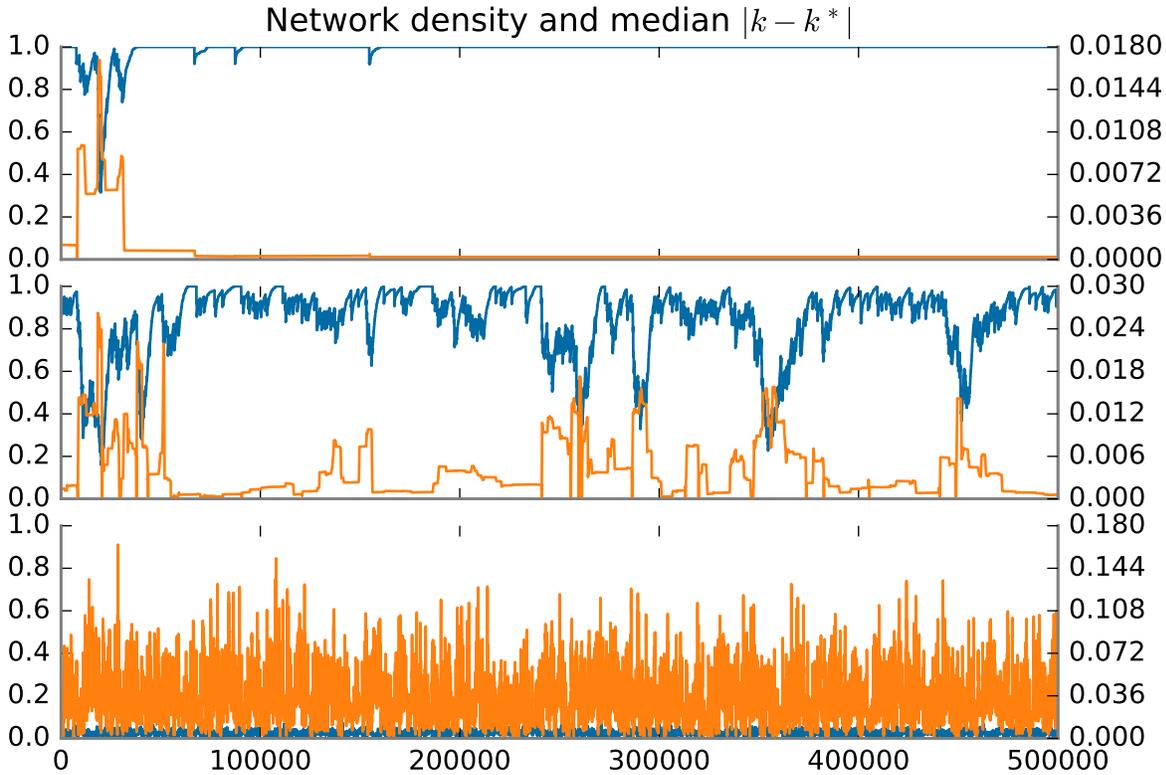


Figure 1: Principal types of dynamics in the model.

Time series of network density and median agent age for three runs of the model. The center graph depicts the model with parameters as in table 1. The models for the top and bottom graphs differ only in β_0 , with $\beta_0 = 0$ in the top graph and $\beta_0 = 2 \times 10^{-3}$ in the bottom. Network density (blue), i.e. total connections divided by $N(N - 1)$, is on the left axis; median $|k - k^*|$ (orange), i.e. median distance from k^* , is on the right axis.

⁵The model is written in the Python programming language (Python Software Foundation) using the NetworkX (Hagberg et al. (2008)) and Pandas (McKinney (2012)) packages. Seed values for runs are randomly chosen but are stored for potential later use. We establish a “burn-in” period at the beginning of a run of the model. At the beginning of the burn-in period, all agents have uniformly drawn k -values and zero connections. The burn-in period ends when all agents have died at least once, after which the model run is considered to begin and time is reset to $t = 0$.

Figure 1 depicts time series of the number of connections among agents and the median distance (across agents) of opinions from the state of the world for three parameterizations of the model. The baseline version of the model is the center graph. The model parameters used in the top and bottom graphs differ only in the exogenous probability of death, β_0 . The top graph depicts the model with $\beta_0 = 0$ and the bottom graph depicts the model with $\beta_0 = 2 \times 10^{-3}$, so that β_0 increases from the top graph to the bottom graph. Each figure graphs the *network density*, defined as the total connections as a proportion of the maximal possible total connections (which is equal to $N(N - 1)$), in blue (left axis) and the median value across agents of the distance of opinions from the state of the world, that is, the median $|k - k^*|$, in orange (right axis). The model is run over 500,000 periods. The time axis is common to all three graphs.

Figure 1 reveals the three principal types of dynamics that occur in the model. Consider the top figure, with zero exogenous probability of death. At most times during the model run the network density is nearly one, indicating that at these times the network is nearly fully connected. Moreover, the distance of the median opinion to the state of the world is very close to zero, suggesting that agents have achieved consensus and that their opinions are very accurate. With the exception of the first 50,000 periods, the effect of the few deaths that occur is idiosyncratic, as connections are rebuilt and information sharing quickly resumes.

Contrast the top and bottom graphs in figure 1. In the bottom graph, the exogenous probability of death is higher (equal to $\beta_0 = 2 \times 10^{-3}$), and the resulting model has a completely different dynamic from the top graph. Network density is uniformly below 0.1, indicating that at all times the agents are only sparsely connected. Moreover, the median distance of opinions from the state of the world is highly volatile with an average of about 0.05. Connections never reach a significant level as the rate of death due to exogenous reasons overwhelms the link formation process, and thus impedes information sharing. This then feeds back into an increased probability of agent death.

The dynamics of the baseline model in the center graph of figure 1 are more interesting. Looking at the time series of the network density and opinions, one observes pronounced cycles in which periods of high connectivity and relatively accurate opinions transition quickly into periods of low connectivity and relatively inaccurate opinions (crashes), eventually transitioning again to a period of high connectivity and relatively more accurate opinions (rebuilding). These cycles are stochastic in the sense that they occur irregularly and last for varying amounts of time. Finally, referring again to the top graph where $\beta_0 = 0$, note that from the first 50,000 periods in the top graph that a cycle occurs even in the case when the exogenous probability of death is zero.

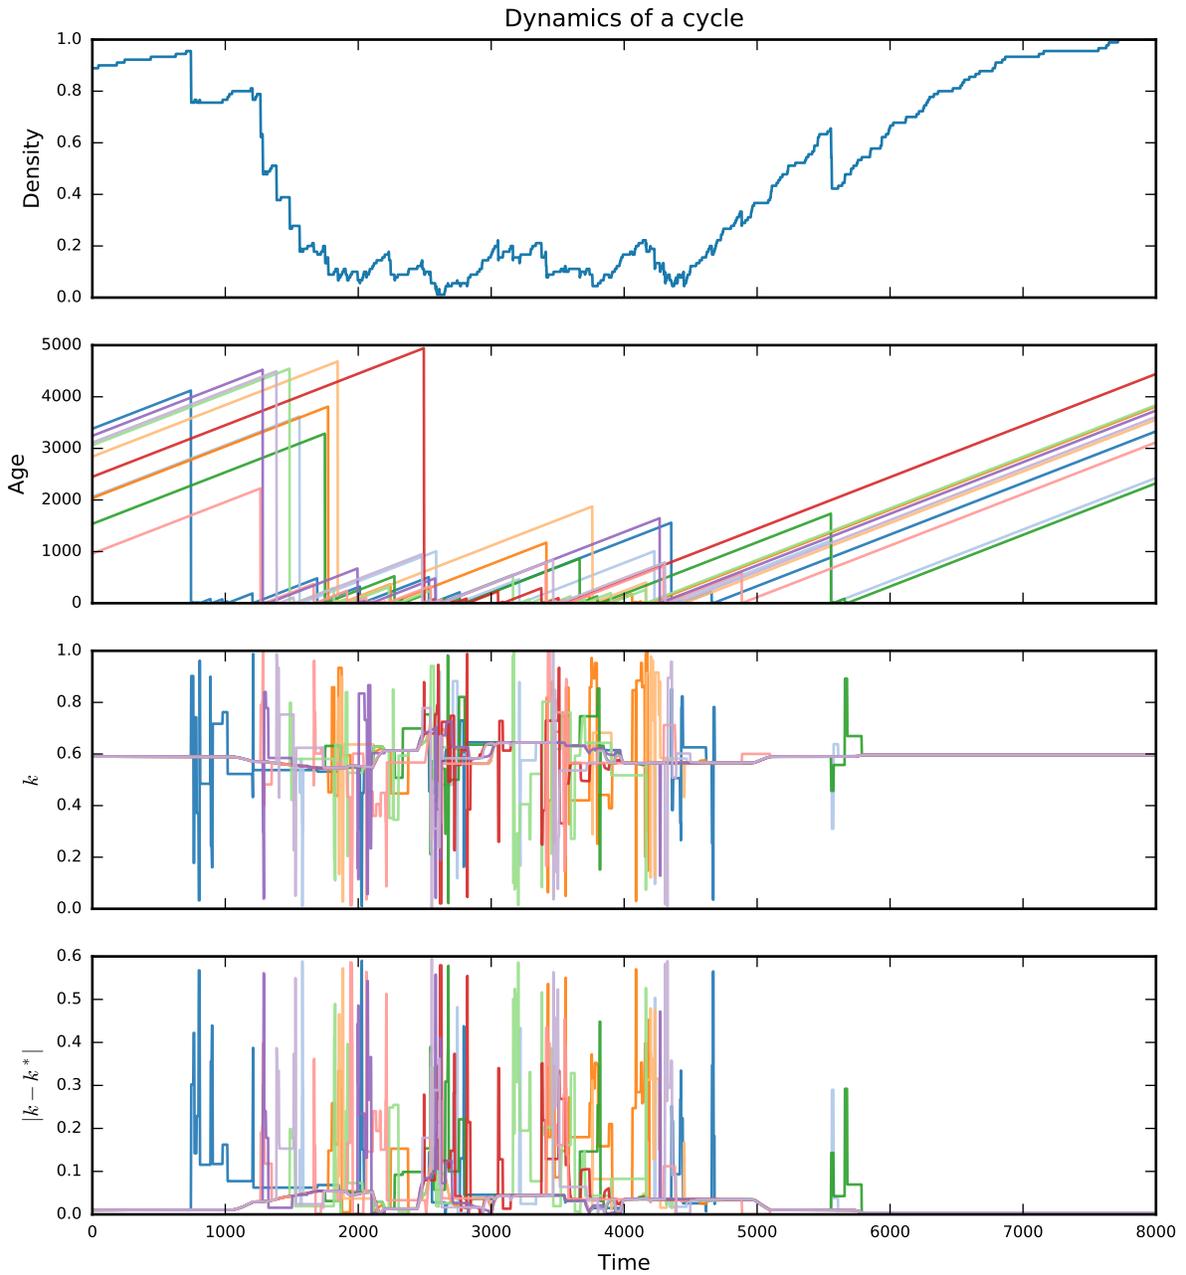


Figure 2: Detail on dynamics of a typical cycle.

Time series for an 8,000 period cycle are depicted from a simulation of the model under the baseline parameters of table 1, except that $N = 10$. The top graph depicts network density. The bottom three graphs depict attributes for all of the agents, which are represented by the same colors in the graphs. The second graph depicts agent age, the third graph depicts agent opinions (k -values), and the bottom graph depicts the distance of opinions from the state of the world ($|k_i - k^*|$). The time axis is common to all graphs.

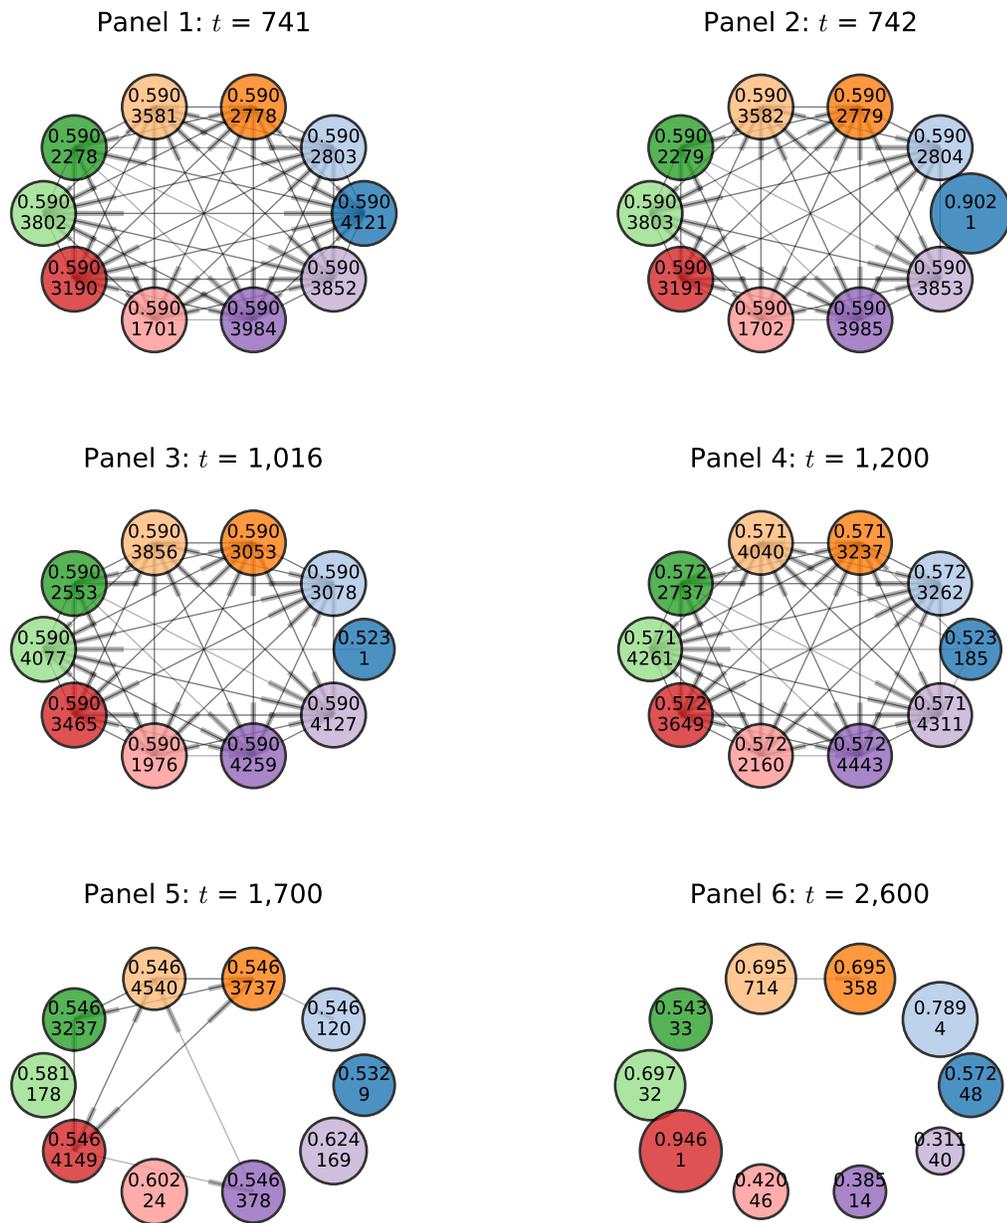


Figure 3: Crash: Transition from high to low connectivity.

Snapshots of the graph representations of the network at specified times throughout the cycle depicted in figure 1. Agents are represented by circular nodes with the same colors as in the bottom three graphs of figure 1. Connections among agents are represented by edges, with the dark gray end indicating the direction of information flow. Nodes are labeled with the corresponding agent's opinion and age.

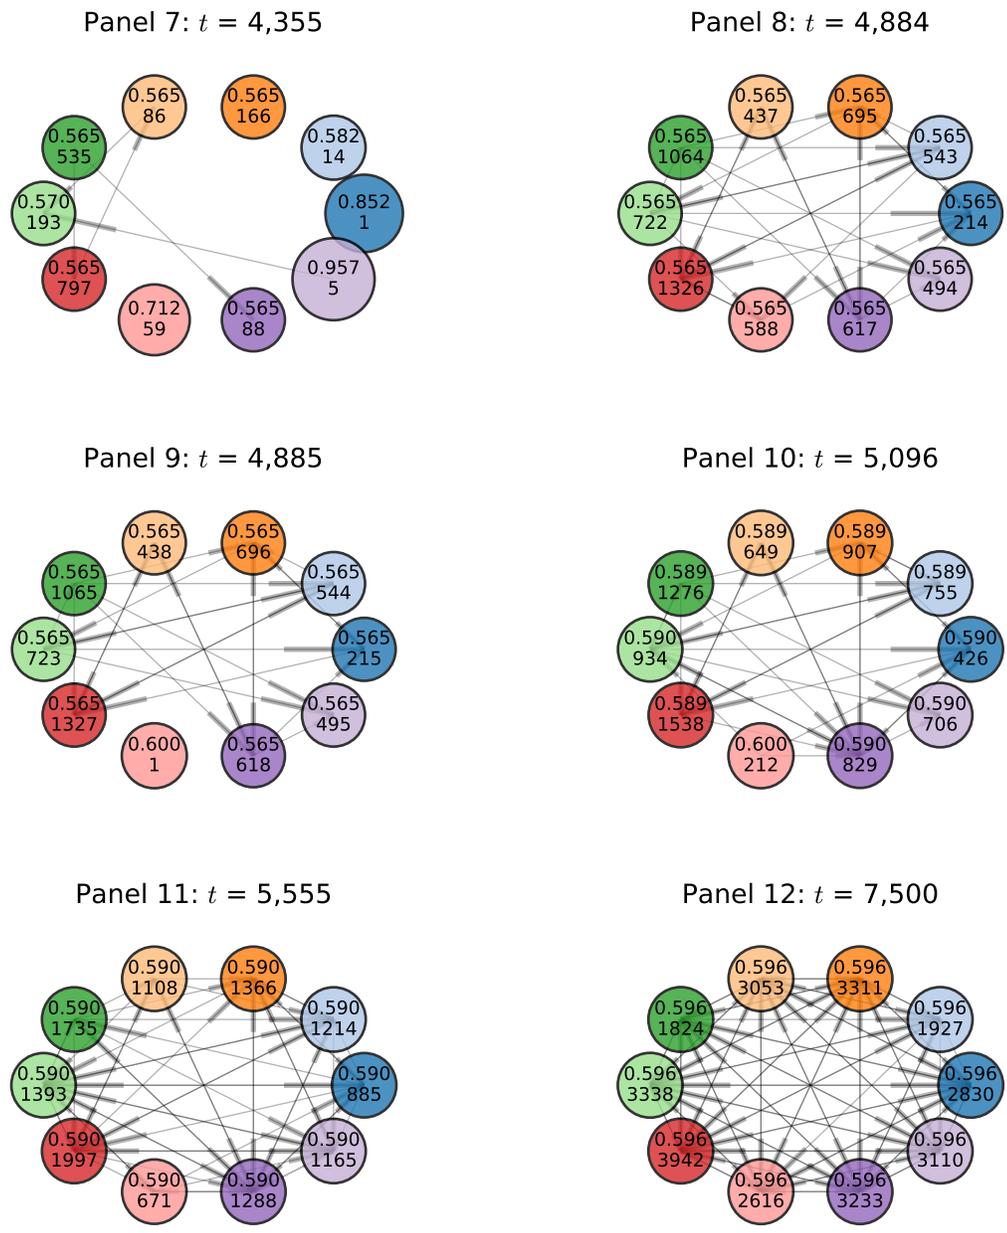


Figure 4: Rebuilding: Transition from low to high connectivity.

Snapshots of the graph representations of the network at specified times throughout the cycle depicted in figure 1. Agents are represented by circular nodes with the same colors as in the bottom three graphs of figure 1. Connections among agents are represented by edges, with the dark gray end indicating the direction of information flow. Nodes are labeled with the corresponding agent's opinion and age.

Figures 2, 3 and 4 aid in explaining the cycles. Figure 2 depicts the detailed dynamics of a cycle from a model run with parameters as in table 1, except that the number of agents is reduced to $N = 10$. It contains four time series graphs, where the time axis is shared among all four. The top graph depicts the network density, which ranges from zero to one, indicating that the network evolves from totally disconnected to a complete graph over the course of the cycle. The remaining three graphs depict the ages, opinions, and distances of opinions from the state of the world, respectively, for each of the 10 agents. Recall that, in any given period, the age of an agent that has survived the previous period is incremented by one, whereas an agent that died is replaced with a new agent with age equal to zero. We depict by the same color the agent that replaces a dead agent, so that if a blue agent dies then its replacement is also blue. For the ease of exposition we will call each color a *color seat*. Thus at any given time each agent is identified to its color seat, and when an agent dies its replacement occupies the same seat.

Figures 3 and 4 provide snapshots of the network at selected times during the cycle depicted in figure 2. The figures are composed of 12 panels ordered in time, where the specific time depicted is indicated in the panel's title. In each panel the agents are represented by circular nodes, with colors corresponding to the colors of the agents in the lower three graphs of figure 2. The nodes in each panel are labeled with the corresponding agent's opinion (rounded to three decimal places) and age at that time, and the size of the node is proportional to the agent's opinion. Connections between agents are represented by the edges of the graph in the panel. The thick end of an edge, colored in dark gray, indicates the direction of information flow; an edge whose ends are both thick represent bidirectional information flow (these edges count as two edges in the graph of network density in figure 2).

The cycle depicted in figures 2, 3 and 4 occurs over 8,000 periods. From time $t=0$ to $t=741$ the network is in a highly connected state, with network density above 0.9, and agents are in agreement on their opinions, with k -values close to the state of the world k^* . This is confirmed in panel 1 of figure 3, which shows the graph representation at time $t=741$ and from which one observes agents in a highly connected state having achieved consensus opinions of 0.590, close to the state of the world 0.6. Moreover, the second graph of figure 2 indicates that at time $t=741$ agent ages are dispersed, ranging from 1,701 to 4,121 periods. Subsequently, at time 742, the blue agent, which is the oldest agent, dies, either because its opinion of 0.590 is not precisely equal to the state of the world, or as a result of the exogenous probability of death being greater than zero. Panel 2 of figure 3 shows that the blue agent's death results in the destruction of all of its connections and that a new agent in its color seat is born at $t=742$ with initial

opinion equal to 0.902. In the graph of ages of figure 2, the blue agent's death at time $t=742$ can be determined by the vertical drop in its age from about 4,121 to zero. Agent deaths throughout the cycle can be determined in the same way.

Between times $t=742$ and $t=1,015$, all agent deaths occur in the blue seat. The agents replacing dead agents are born into the blue seat with randomly drawn opinions that are generally different from the other agents and the state of the world. They struggle to connect and often die. Then at time $t=1,016$, an agent in the blue seat enters the system with k -value approximately equal to 0.523, closer to the state of the world than his recent predecessors in the blue seat, but still different from the other agents. Partly because of its relatively more accurate initial opinion, this agent survives long enough for at least some of the other agents to connect to it and obtain information from it. Importantly, the agent spreads its relatively inaccurate opinions while not forming in-links and obtaining more accurate information from the rest of the agents. This occurs (in expectation) because the blue agent connects to another agent with probability p^c in a period, but each of the other agents, having already connected to each other, have that same probability of connecting. For instance, panel 3 of figure 3 shows that the light green agent has already connected to the just-born blue agent. Panel 4 of figure 3, 184 periods after panel 3, shows that by time $t=1,200$ four agents (orange, light orange, light green, and light purple) have connected to the blue agent, but the blue agent has no in-links itself and is not receiving information from any of the other agents. Its generally inaccurate information spreads to the other agents and, although the blue agent has relatively little weight in each of the other agents' update rule, the fact that its poor information is accounted for in at least some of the other agents' updating and that the other agents all update according to each others' information means that the blue agent's poor information has outsize influence as it spreads nonlinearly to the other agents. As we see from period $t=1,016$ to $t=1,200$, the k -values of *all* of the other agents deteriorate in accuracy from 0.59 to 0.573, which is closer to the k -value of the blue agent.

To facilitate understanding of this dynamic, consider the following simplified example. Assume there are $N = n + 1$ agents. The first n agents are connected to each other and the last agent is connected to the second-last agent in a one-way transmission of information wherein the second-last agent receives information from the last agent but not vice-versa. The last agent thus does not receive information from any of the other agents; it represents perhaps a new entrant in the system, such as the blue agent at time $t=742$ above. Suppose the first n agents have opinions exactly equal to the state of the world $k^*=0.6$ but the last agent has opinion equal to 0.5. For simplicity, assume the first $n - 1$ agents equally weight their connections' opinions, the second-last agent

places a small weight of $\epsilon > 0$ on the last agent's opinions and equally weights the others' opinions, and the last agent, the new entrant, only considers his own opinion (cf. panels 3 and 4 of figure 3). Also assume that there is no link formation or destruction process. This is an example of a static DeGroot learning model. The weight matrix \mathbf{W} of (1) is constant in time, and is equal to

$$\mathbf{W} = \begin{bmatrix} \frac{1}{n} & \frac{1}{n} & \dots & \frac{1}{n} & 0 \\ \frac{1}{n} & \frac{1}{n} & \dots & \frac{1}{n} & 0 \\ \dots & \dots & \dots & \dots & \dots \\ \frac{1-\epsilon}{n} & \frac{1-\epsilon}{n} & \dots & \frac{1-\epsilon}{n} & \epsilon \\ 0 & 0 & \dots & 0 & 1 \end{bmatrix}. \quad (9)$$

The agents' opinions t periods in the future are given by the computation

$$\mathbf{k}(t) = \mathbf{W}^t \mathbf{k}(0) = \begin{bmatrix} \frac{1}{n} & \frac{1}{n} & \dots & \frac{1}{n} & 0 \\ \frac{1}{n} & \frac{1}{n} & \dots & \frac{1}{n} & 0 \\ \dots & \dots & \dots & \dots & \dots \\ \frac{1-\epsilon}{n} & \frac{1-\epsilon}{n} & \dots & \frac{1-\epsilon}{n} & \epsilon \\ 0 & 0 & \dots & 0 & 1 \end{bmatrix}^t \begin{bmatrix} 0.6 \\ 0.6 \\ \dots \\ 0.6 \\ 0.5 \end{bmatrix}.$$

Using known results from Golub and Jackson (2010), we deduce that the agents in this simple example converge to a consensus opinion equal to the opinion of the last agent.⁶ That is,

$$\lim_{t \rightarrow \infty} \mathbf{k}(t) = \begin{bmatrix} 0.5 \\ 0.5 \\ \dots \\ 0.5 \\ 0.5 \end{bmatrix}. \quad (10)$$

This is perhaps surprising given the isolation of the last agent relative to the rest of the highly interconnected agents and the arbitrarily small influence that the last agent's opinion has on merely one of the other agents. Moreover, the exponential rate at which information diffuses among the first n highly interconnected agents, relative to the small probability that the last agent connects into the highly interconnected group, implies that the above example, while simplified, nonetheless captures an essential point about our model that a single agent with an inaccurate opinion can "infect" the opinions of all the other agents, even if only one other agent is directly connected to it.

⁶Using the terminology in Golub and Jackson (2010), it is easily seen that the singleton set containing the last agent is the only so-called closed group, and thus the only minimal closed group, relative to \mathbf{W} in (9). Then (10) follows from Theorem 3 therein.

The blue agent's infection of the other agents sets the stage for the crash, the disorderly transition from a state of high connectivity and accurate consensus to a state of low connectivity and widely dispersed opinions. Five agents die, including three of the oldest, between $t=1,200$ and $t=1,564$ partly because of their less accurate opinions due to the influence of the blue agent. Their replacements, with randomly drawn beliefs, have generally poor information and either die as they struggle to connect, or stay in the system long enough to spread their poor information. The bottom graph in figure 2 indicates that opinions continue to get farther from the state of the world over the period from about $t=1,600$ to $t=2,100$, precipitating the deaths of the remaining old agents and leading eventually to almost a complete collapse of the links in the network.

The ensuing period between times $t=2,100$ and $t=4,500$ is one of low connectivity and high dispersion in opinions. For example, panel 6 in figure 3 shows that at time $t=2,600$ most of the agents are young and have inaccurate opinions, and that there is only one connection among agents at that time. During the low-connectivity state, agent deaths are common since they are unable to obtain a sufficient number of connections to improve their opinions. There are a few periods when opinion volatility settles and connections rebuild somewhat so that it seems that the system might transition back to a high connectivity state, but this proves temporary and the system reverts to the low-connectivity state. Finally, around $t=4,400$, half of the agents in the system die in close succession and are replaced eventually with agents whose opinions are close enough to the state of the world to remain in the system long enough to connect and obtain more accurate opinions. The network begins to rebuild as density goes from about 0.1 to 0.3 between $t=4,500$ and $t=4,800$, its highest value since the crash. Panel 8 in figure 4 shows that by time $t=4,884$ agents are moderately connected and have achieved a consensus opinion around 0.565, a distance of 0.035 away from the state of the world. Then, as panel 9 shows, at time $t=4,885$, the pink agent dies and is replaced with an agent whose initial opinion is equal to 0.60029, which is very close to the state of the world. Since its opinion is so close to the state of the world, the pink agent remains in the system long enough to spread its accurate opinions. There are no deaths in the period between $t=4,885$ and $t=5,096$ and panels 9 and 10 in figure 4 show that the three out-links to which the pink agent is a party have resulted in the spread of its accurate information to all of the rest of the agents, as their consensus opinion has improved from 0.565 to 0.589. Thereafter, agents are in agreement about opinions and these opinions are close to the state of the world. A few idiosyncratic deaths occur as connections increase until the network is fully connected around $t=8,000$, at which time the agents are in consensus and have opinions that are close to the state of the

world (see panels 11 and 12 in figure 4).

To summarize, certain parameterizations produce cyclic dynamics. Cycle frequency depends on the relationship between the probability of connection and the exogenous probability of death (we elaborate on this in section 6). In the high-connectivity state agents are very interconnected and generally have accurate opinions. But because of either the positive exogenous probability of death or slight inaccuracies in their opinions, eventually some agents will die. If a relatively old and interconnected agent with accurate opinions dies, then its failure removes its information from all of its out-links. It is replaced by an agent with generally inaccurate opinions. But if this agent's opinions are sufficiently accurate such that it remains in the system long enough for other agents to connect to it, then it can infect the other agents with its poor information. Despite its possibly low weight in other agents' updating functions, this infection can spread nonlinearly because the agents are interconnected and repeatedly update their opinions. As a result, most or all of the agents' opinions can deteriorate, which can lead to more deaths and the destruction of connections. If the subsequent replacement agents have poor information but are able to connect and spread it to others, then this can lead to a further deterioration in information and to more deaths. Eventually, most connections are severed and agents enter a low-connectivity state in which their opinions are dispersed. Deaths occur frequently in this state. Eventually, agents with sufficiently accurate information survive long enough to share their information with others. Their good information can spread to the other agents in the same way that bad information spread before. Agents with good information live longer and longer, connecting more and more often. Eventually the system enters a period of high connectivity and achieves a consensus at accurate opinions.

6 Comparative statics

Next we present comparative statics on the model's two main sources of randomness, the per-period probability of connection p^c and common exogenous probability of death β_0 in (3). These model parameters represent the speeds of information transmission and destruction, respectively. A higher probability of per-period connection p^c produces a faster rate of information transmission, as agents are able to make connections more often and transmit information over these connections. Conversely, a higher β_0 represents a higher rate of death from exogenous reasons, and thereby a higher rate of information destruction. By examining properties of the model as these parameters vary, we are able to illuminate this trade-off and determine which sets of parameters produce more stability and which produce more fragility.

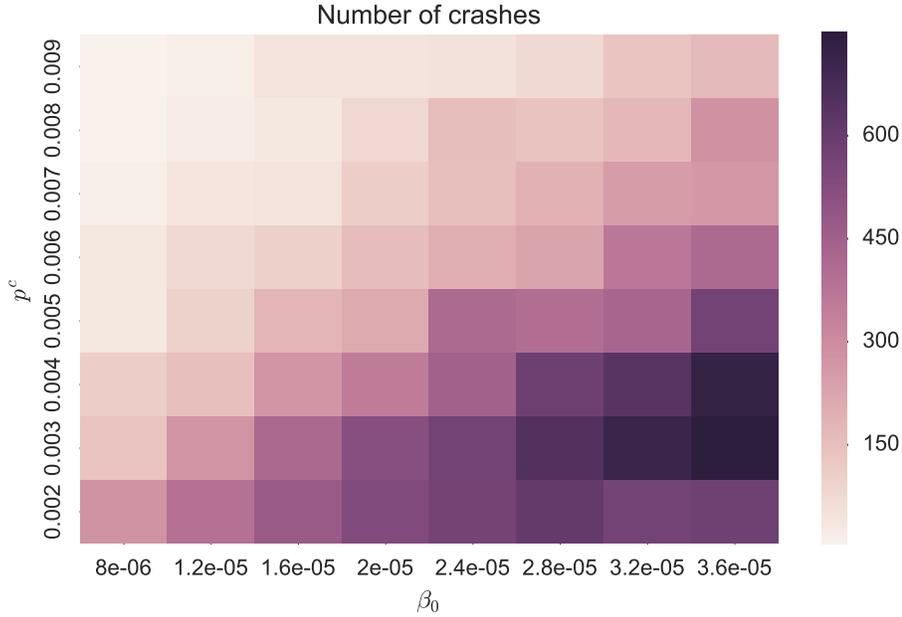


Figure 5: Comparative statics: Crash frequency.

Heat map showing how crash frequency varies as probability of connection (p^c) and exogenous probability of death (β_0) vary. Each (p^c, β_0) -pair is given by the values along the axes and all other model parameters are as in table 1. Each set of parameters is simulated over 25,000,000 periods.

Figure 5 is a heat map that illustrates the relationship between p^c and β_0 with respect to the frequency of crashes. For each (p^c, β_0) -pair indicated along the axes, the model is run over a period of 25,000,000 periods, with all other parameters as in Table 1. Crash frequency is determined as follows. A crash is composed of a transition from a high-connectivity state in which agents have reached an accurate consensus opinion close to the state of the world to a low-connectivity state in which agents' opinions are widely dispersed. A crash is eventually followed by a rebuilding period, which is a transition back to the high-connectivity state. We observe from figures 1 and 2 that network density closely tracks periods of consensus and dispersion of opinions. We will use network density to define the states and compute the crashes because other natural measures of crash count, such as using variance of opinions, tend to be noisier than network density. We consider the system to be in a high-connectivity state when the network density is greater than 0.6, and to be in a low-connectivity state when the network density is less than 0.4. The system is considered to be in an intermediate state when network density is between 0.4 and 0.6.⁷ Thus a crash has occurred if network density goes from above 0.6 to below 0.4. For example, applying this definition to the graphs of network density in figure 1, we count one crash in the top graph, six in the

⁷Other cutoff values for state classification produce similar results.

center graph, and zero in the bottom graph.

We see in figure 5 that for a fixed probability of connection p^c , a higher β_0 generally leads to more crashes. For a fixed β_0 , a lower probability of connection results in more crashes. The greatest number of crashes occurs when the probability of connection is low and the exogenous probability of death β_0 high, and the least number of crashes occurs when connection probability is high and exogenous probability of death is low. Moreover, the level curves for crash frequency appear to be approximately linear in p^c and β_0 . Intuitively, as β_0 rises, the relative value of being close to the state of the world k^* declines. Agents die more arbitrarily, and fail to establish and maintain many centralized connections. This results in a more complete network, with more frequent crashes. However, as β_0 goes to zero, agents that are lucky enough have opinions k close to k^* and become more central because they do not arbitrarily die and agents connecting to them maintain their connections. This is reinforced if p^c is high, and leads to a system with many connections to central nodes. Such a system experiences fewer crashes because central nodes with high-quality information are unlikely to die; however, when a central node dies, the effects of its death propagate to its many connections, leading to a crash.

7 Tribal preference and network stability

A widespread finding in the sociology and social networks literature is that agents associate more with other agents that are similar to them and less with those that are not. This tendency of agents to disproportionately link to similar agents is called homophily. Homophily has been observed across many agent characteristics, including age, gender, race, education level, profession, and political affiliation (see McPherson et al. (2001) for a review). It is natural to consider homophily in our context. In the application of the model to leverage (see section 3), for example, agents can be partitioned into dealer banks and non-dealer entities. Dealer banks are accustomed to doing high volumes of business with each other relative to other financial institutions and thus might be more likely to associate with each other as opposed to non-dealers.

Since homophily affects the propensity of agents to connect to other agents, it has a direct effect on the link-formation process and on the flow of information in the network. This effect on information flow was analyzed by Golub and Jackson (2012). In a model incorporating homophily they find that, for average-based updating processes such as the DeGroot updating method, the time to converge to consensus is increasing and convex in homophily, holding other factors fixed. A natural question in our context is how homophily affects the frequency of crashes.

We analyze the effect of homophily on network stability in a simple extension of our model in which we introduce the notions of a tribe and tribal preference. The *tribe* is a pre-specified subset of the agents of size $T \leq N$ in which the agents are alike in some way that affects their propensity to connect with one another rather than to the general populace. In the application of the model to leverage, the tribe could be a set of dealers, as discussed above.

For members of the tribe, the selection of an agent with whom to connect is made according to the *tribal preference* parameter $p^{tp} \in (0, 1)$. Conditional on connecting in a given period, the tribal preference parameter p^{tp} is defined to be the probability of connecting to a member of the tribe; if an agent is already connected to all other members of the tribe, then it connects randomly to an available agent from the broader population. For simplicity, we assume that tribal preference is the same for all members of the tribe, and that the tribe and tribal preference parameters are fixed at initial time and are not time-varying. In our example, dealers need not transact or learn information from other dealers only. Tribal preference quantifies how willing they are to rely on non-dealers' information.

As the tribal preference parameter p^{tr} increases, agents increasingly prefer to connect to agents who are within their own tribe rather than connect to random agents from the greater population. By restricting the set of possible agents with whom to connect, they are willfully ignoring information because they are less confident of the reliability of information outside their tribe. They are more discriminatory in their connections, and less trusting of outsiders. Conversely, as the tribal preference parameter p^{tr} decreases, agents increasingly prefer to connect to random agents not necessarily within their tribe, expressing a willingness to obtain information from a broad-based supply of agents. Agents are more trusting of other agents, as they are more confident of the quality of information available broadly, and a more global system arises.

Note that the proof of theorem 4.1 goes through for all $p^{tp} \in [0, 1)$ and thus the model with tribal preference is ergodic. We simulate the model for various tribe and population sizes, with other parameters as in table 1. Figures 6 and 7 are representative of the resulting behavior. Figure 6 shows that the connectivity of the system depends on how strongly agents prefer to connect to agents of their type. The figure depicts the empirical frequency distributions of total connections for various levels of tribal preference ranging from zero to one, as indicated in the graph's legend. The tribe size is $T = 25$, the population size is $N = 100$, and the model is simulated over 10,000,000 periods. At low levels of tribal preference, there are fewer interconnections among tribe members, as the tribe connects about as willingly to the greater population of agents. As tribal preference increases, mass moves from low numbers of within-tribe

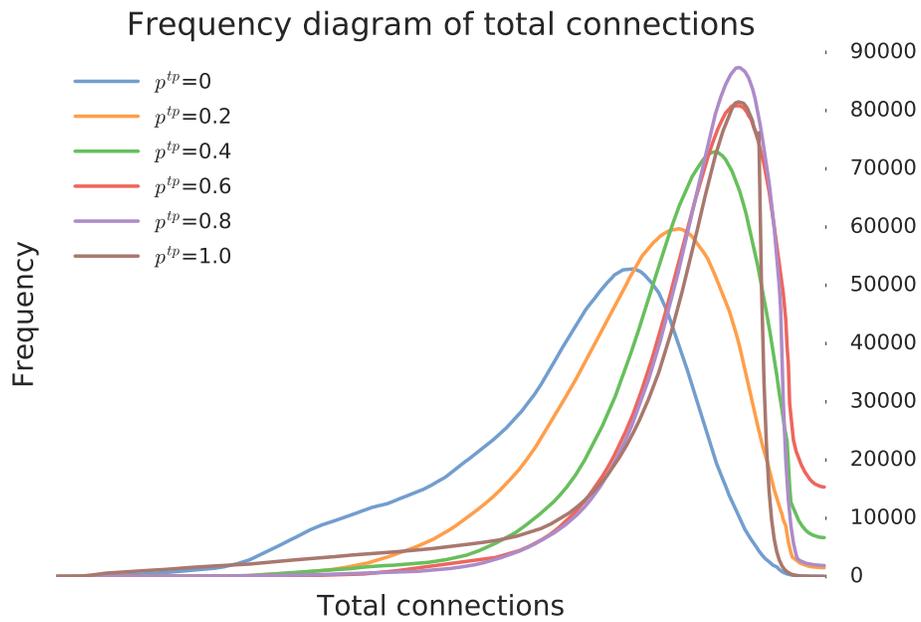


Figure 6: Density of within-tribe connections.

Depiction of the frequency distributions of total within-tribe connections for various levels of tribal preference indicated in the graph legend. The tribe size is 25 and the population size is 100. All other parameters of the model are as in table 1. Graph is based on model runs of 10,000,000 periods. The data are smoothed using a Savitzky-Golay low-pass filter (see Savitzky and Golay (1964)).

connections to high numbers of within-tribe connections. This result confirms that the tribal preference parameter governs the interconnections among tribe members.

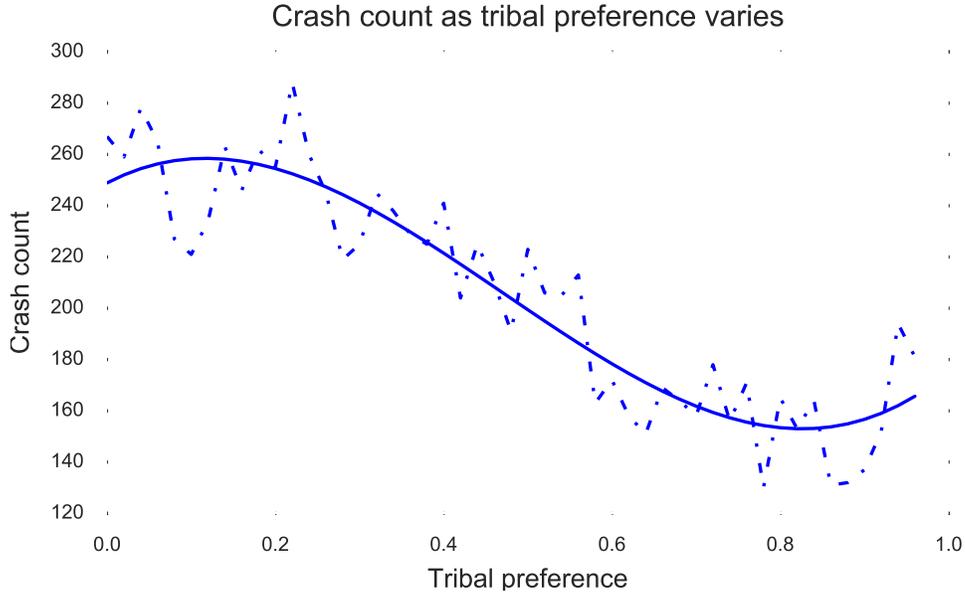


Figure 7: Within-tribe crash count as tribal preference varies.

Depiction of the crash frequency as tribal preference varies in $[0, 1]$. The tribe size is 12 and the population size is 48. All other parameters of the model are as in table 1. Results are based on model runs of 25,000,000 periods. The dotted line depicts the raw data while the continuous line depicts the raw data fitted by a cubic polynomial.

We next consider the relationship between tribal preference and stability of the system. We summarize this relationship in terms of the frequency of crashes. Figure 7 shows how the dynamics within the tribe are affected by the tribal preference parameter. The tribe size is $T = 12$, the population size is $N = 48$, and the model is run over 25,000,000 periods. The figure depicts, for varying levels of tribal preference ranging from zero to one, the number of crashes that occur in the model simulation. A crash is composed of a transition from a period of high connectivity and consensus among agents to a period of low connectivity and dispersed opinions. As in section 6, this is operationalized using network density cutoffs of 0.6 and 0.4. Finally, the solid blue is determined by fitting a cubic polynomial to the raw data.

Figure 7 shows that tribal preference determines crash frequency. As tribal preference increases, i.e. agents become less trusting of outsiders, the number of crashes decreases from about 250 to 150. For low levels of tribal preference, all agents in the tribe are connecting indiscriminately among the larger population. The agents are thus exposed, directly and indirectly, to a greater number of agents. The more agents there are, the higher the chance that at least one agent dies (either for exogenous reasons or

because of imperfect opinions). As discussed in section 5, if a dead agent is replaced by a new agent with poor information relative to the other agents but good enough information in an absolute sense to survive and connect, then in a highly connected state this relatively poor information can propagate widely, reducing the accuracy of others' information, and leading to more deaths and eventually a crash. Conversely, for high levels of tribal preference all agents within the tribe are mainly connecting to themselves. As the tribe is a proper subset of the general population, this leads to fewer chances for a dead agent's replacement to infect the tribe and cause a crash. Therefore, over periods of the same length, lower tribal preference leads to a greater number of crashes.

8 Conclusion

We develop a social networks model that combines elements of DeGroot learning and dynamic network formation in which continued membership in the network is contingent on the accuracy of opinions. In contrast to results in the literature on DeGroot learning, we find through simulations that the model cycles between a state of high connectivity and consensus and a state of low connectivity and dispersed opinions. Transitioning from a state of high to low connectivity occurs when a highly connected agent dies and is replaced by an unconnected agent with relatively poor quality information, which then spreads through the highly interconnected network, leading to less accurate opinions and to more deaths. The transition from a low to high state of connectivity occurs when agents with accurate opinions stay alive long enough to spread their information as connections form. A comparative statics analysis on the parameters governing the cycles suggests that level curves of crash frequency are approximately linear in the probability of connection and the exogenous probability of death. In an extension of the model that incorporates homophily, we find that the frequency of crashes is decreasing in the degree of homophily.

The model shows that interesting macroscopic phenomena can emerge as the result of simple behavioral rules at the agent level. The complexity resulting from the nonlinear, state-dependent dynamics makes the model largely intractable, though it is still ergodic.

References

- D. Acemoglu and A. Ozdaglar. Opinion dynamics and learning in social networks. *Dynamic Games and Applications*, 1(1):3–49, March 2011.
- V. Bala and S. Goyal. A noncooperative model of network formation. *Econometrica*, 68(5): 1181–1229, 2000.
- A. G. Chandrasekhar, H. Larreguy, and J. P. Xandri. Testing models of social learning on networks: Evidence from a lab experiment in the field. NBER Working Paper No. 21468, National Bureau of Economic Research, August 2015.
- M. H. DeGroot. Reaching a consensus. *Journal of the American Statistical Association*, 69(345): 118–121, 1974.
- P. M. DeMarzo, D. Vayanos, and J. Zwiebel. Persuasion bias, social influence, and unidimensional opinions. *Quarterly Journal of Economics*, 118(3):909–968, 2003.
- J. Geanakoplos. Promises, promises. In W. B. Arthur, S. Durlauf, and D. Lane, editors, *The Economy as an Evolving Complex System*, volume 2, pages 285–320. Addison-Wesley, 1997.
- J. Geanakoplos. Liquidity, default, and crashes: Endogenous contracts in general equilibrium. In M. Dewatripont, L. P. Hansen, and S. J. Turnovsky, editors, *Advances in Economics and Econometrics: Theory and Applications*, volume 2 of *Econometric Society Monographs: Eighth World Congress*, pages 170–205. Cambridge University Press, 2003.
- J. Geanakoplos. The leverage cycle. In D. Acemoglu, K. Rogoff, and M. Woodford, editors, *NBER Macroeconomics Annual 2009*, volume 24, pages 1–65. University of Chicago Press, 2010.
- B. Golub and M. O. Jackson. Naive learning in social networks and the wisdom of crowds. *American Economic Journal: Microeconomics*, 2(1):112–149, 2010.
- B. Golub and M. O. Jackson. How homophily affects the speed of learning and best-response dynamics. *The Quarterly Journal of Economics*, 127(3):1287–1338, 2012.
- V. Grimm and F. Mengel. An experiment on belief formation in networks. Available at SSRN: <http://ssrn.com/abstract=2361007>, July 2014.
- A. A. Hagberg, D. A. Schult, and P. J. Swart. Exploring network structure, dynamics, and function using NetworkX. In *Proceedings of the 7th Python in Science Conference (SciPy2008)*, pages 11–15, Pasadena, CA USA, Aug. 2008.

- R. Hegselmann and U. Krause. Opinion dynamics and bounded confidence, models, analysis and simulation. *Journal of Artificial Societies and Social Simulation*, 5(3), 2002.
- R. Hegselmann and U. Krause. Opinion dynamics driven by various ways of averaging. *Computational Economics*, 25(4):381–405, 2005.
- O. C. Ibe. *Markov Processes for Stochastic Modeling*. Elsevier Academic Press, 2009.
- M. O. Jackson and B. W. Rogers. Meeting strangers and friends of friends: How random are social networks? *The American Economic Review*, 97(3):890–915, 2007.
- M. O. Jackson and A. Watts. The evolution of social and economic networks. *Journal of Economic Theory*, 106:265–295, 2002.
- O. Kallenberg. *Foundations of Modern Probability*. Springer, Second edition, 2002.
- J. Lorenz. A stabilization theorem for dynamics of continuous opinions. *Physica A: Statistical Mechanics and its Applications*, 355(1):217–223, 2005.
- J. Lorenz. Convergence of products of stochastic matrices with positive diagonals and the opinion dynamics background. *Positive Systems*, pages 209–216, 2006a.
- J. Lorenz. Consensus strikes back in the Hegselmann-Krause model of continuous opinion dynamics under bounded confidence. *Journal of Artificial Societies and Social Simulation*, 9(1), 2006b.
- P. V. Marsden. Core discussion networks of Americans. *American Sociological Review*, pages 122–131, 1987.
- P. V. Marsden. Homogeneity in confiding relations. *Social Networks*, 10(1):57–76, 1988.
- W. McKinney. *Python for Data Analysis: Data Wrangling with Pandas, NumPy, and IPython*. O’Reilly, 2012.
- M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, pages 415–444, 2001.
- A. Mirtabatabaei and F. Bullo. Opinion dynamics in heterogeneous networks: Convergence conjectures and theorems. *SIAM Journal on Control and Optimization*, 50(5):2763–2785, 2012.
- M. Mueller-Frank and C. Neri. Social learning in networks: Theory and experiments. Available at SSRN: <http://ssrn.com/abstract=2328281>, 2013.

- E. Pardoux. *Markov Processes and Applications: Algorithms, Networks, Genome and Finance*. John Wiley and Sons, 2008.
- Python Software Foundation. Python language reference, version 2.7. Available at <http://www.python.org>.
- A. Savitzky and M. J. E. Golay. Smoothing and differentiation of data by simplified least squares procedures. *Analytical Chemistry*, 36(8):1627–1639, 1964.
- B. Skyrms and R. Pemantle. A dynamic model of social network formation. In *Adaptive Networks*, pages 231–251. Springer, 2009.
- M. Staudigl. Co-evolutionary dynamics and Bayesian interaction games. *International Journal of Game Theory*, 42(1):179–210, 2013.
- A. Watts. A dynamic model of network formation. *Games and Economic Behavior*, 34(2):331–341, 2001.
- G. Weisbuch, G. Deffuant, F. Amblard, and J. Nadal. Meet, discuss, and segregate! *Complexity*, 7(3):55–63, 2002.