Games and Economic Behavior 150 (2025) 312–333

Contents lists available at ScienceDirect

Games and Economic Behavior

journal homepage: www.elsevier.com/locate/geb

Computing agents' reputation within a network

Federico Grigis [a], Sergio Ortobelli Lozza [b], Sebastiano Vitali [c,d,*]

[a] Risk Department, Banca Popolare di Sondrio, Milan, Italy
[b] Department of Management, University of Bergamo, Via dei Caniana 2, 24127 Bergamo, Italy
[c] Department of Economics, University of Bergamo, Via dei Caniana 2, 24127 Bergamo, Italy
[d] Department of Probability and Mathematical Statistics, Faculty of Mathematics and Physics, Charles University, Sokolovska 83, 186 75 Prague, Czech Republic

ARTICLE INFO

JEL classification:
D83
D85
A14
G11

Keywords:
Social networks
Learning
Social influence
Option structures
Portfolio optimization

ABSTRACT

We propose a model of information transmission and reputation building within a social network that exploits portfolio theory and option structures. The network aims to estimate an unknown parameter through multiple communication rounds. At every communication round, estimates of different agents' abilities are shared, avoiding the repetition of information. These estimates are interpreted as financial assets driven by a compound Poisson process. After every communication round, agents construct a fictitious portfolio of options whose underlying is the vector of shared estimates. The portfolio's weights are exploited to aggregate the information received in the communication round. Sufficient conditions for reaching consensus or polarization are provided.

1. Introduction

Gathering and aggregating information from a variety of information sources is often a crucial step in a decision-making process, especially in the presence of repetition of information. Furthermore, an accurate assessment of the reliability of all available sources is fundamental to making the right decision. To tackle these two issues, this paper proposes two main contributions:

- a novel interpretation of the informative content of the network that excludes the repetition of information over time;
- an innovative procedure that exploits option structures and portfolio theory to determine the weights that agents should assign to their peers within a network.

Indeed, we consider rational and trustworthy agents that form a fully connected network and share, at every communication round, a different piece of information. At the first communication round, agents spread their initial estimates of the unknown parameters; at the second communication round, agents share the weights they assigned to their sources after the first communication round, representing their assessments of agents' abilities to estimate the parameter; at the third round, agents communicate the weights assigned after the second communication round, reflecting their estimates of agents' abilities to evaluate their proficiency in estimating the parameter, and so on. This approach seems reasonable and consistent with intuition. In summary, suppose that after the first communication round, agent i has assigned a high (low) weight to some agent j. Subsequently, agent i may wonder whether other

* Corresponding author.
E-mail address: sebastiano.vitali@unibg.it (S. Vitali).

https://doi.org/10.1016/j.geb.2025.01.002
Received 1 October 2022

Available online 20 January 2025

agents agree with her opinion. Moreover, agent i could be interested in the weight agent j assigned to herself. The lower agent i's confidence in her own assessments, the more valuable this piece of information is. At the same time, agent i could be willing to make other agents aware of the weight she assigned to herself, which reflects her perception of the quality of her own estimates. Note that agents are willing to share their own information truthfully because, by assumption, the network aims to estimate correctly a vector of unknown parameters. Beyond this intuitive argument, our approach is found to reflect real-life interactions, especially in social networks. In a Twitter conversation, or under a Facebook post, the discussants usually first share their idea about a certain topic; then, they begin to support or oppose other participants' opinions. Thus, a second communication round naturally emerges to spread estimates of the reliability of the network's members. Given this new piece of information, every agent assigns to her peers new weights that reflect her perception of the other agents' abilities to evaluate their proficiency at estimating the unknown parameter. Again, agents may be eager to argue about these new assessments in a third round of communication, and so on. Every piece of information exchanged at each communication round contributes to defining sources' overall reliability and eventually the network structure, which is, therefore, endogenously determined. Agents' rationality is thus reflected in their awareness of the whole informative content of the network and in the way they exploit it. A natural consequence of the outlined setting is the relaxation of the assumption of time-constant weights, provided that a different ability is evaluated after each interaction. A procedure to determine the weights to be assigned to each agent after every communication round is then required. Whilst many models propose some evolutionary law driving time-dependent weights, how agents should actually compute their initial weights remains a widely unexplored topic. In DeMarzo et al. (2003), each agent receives a noisy signal about the unknown parameter and then estimates the variance of every peer's error term to assign a weight to her. However, variance estimation can seem like a demanding task, especially when other agents' signals are considered. As mentioned above, we exploit portfolio theory to compute weights. Assume that, within the network, each agent i's payoff is a decreasing function of the distance between the true value of the unknown parameter θ and its estimate. This mimics the payoff of a financial instrument named *butterfly spread*[1] with agent i's estimate of θ as underlying and the θ as center. Thus, assigning a positive weight to one of the available estimates is equivalent to (virtually) allocating a share of wealth to a butterfly with the shared estimate as center. We assume that each agent composes, after the first communication round, a mean-variance optimal portfolio of butterflies, each with a different center corresponding to the estimate of θ by an agent of the network. The vector of portfolio weights computed after the first communication round is shared at the second communication round. The underlying of a butterfly is now agent i's estimate of the vector of weights, and its center is the estimated vector actually shared by an agent j. Again, each agent could choose among different butterflies, each centered on a different vector, and compose a second (virtual) mean-variance optimal portfolio. The vector of weights characterizing this new portfolio is shared at the third communication round and represents the underlying of the butterflies available after the third interaction. A third (virtual) mean-variance optimal portfolio of butterflies is then composed. The same reasoning is applied at every communication round. At first sight, the estimated value of a parameter is quite different from the price of a financial asset that fluctuates over time. However, once the final decision is made, agents may receive new information about either the unknown parameter or the reliability of the agents in their listening set.[2] This could follow, for example, from an enlargement of agents' initial networks. Furthermore, even if the structure of the network were fixed, the reputation of the agents within the network might evolve over time. Therefore, as a stock's price evolves according to the information received by the market, the estimate of an unknown parameter might change over time as agents receive new information. Thus, a butterfly could become more or less profitable because the value of the underlying evolves over time according to the future flow of information. Notably, when applied to practical situations, our model leads to results consistent with intuition (see section 5). Both consensus and polarization of opinions are allowed, depending on agents' self-confidence and their expectations about the future flow of relevant information. The assumption of risk-averse and non-self-confident agents results in a strong preference for diversification and leads, under mild conditions, to unidimensional opinions[3] and, in turn, to consensus. Conversely, agents with high (or rapidly growing) self-confidence would stick to their estimates, causing polarization of the opinions regardless of their risk aversion. Finally, unidimensionality of opinions is obtained regardless of any specific functional form for the listening matrix. This represents, to the best of our knowledge, a minor contribution to the theory of time-inhomogeneous Markov chains.

1.1. Related literature

Our model belongs to the literature on non-Bayesian opinion aggregation processes as it builds upon the crucial assumption that agents update their beliefs through an average-based updating process. This approach can be traced to DeGroot (1974)'s seminal paper, which considers a network with no strategic communication where genuinely new information is spread only at the first communication round. With each following interaction, each agent shares a linear combination of beliefs from the previous communication round. A standard result in this body of literature is that, under mild regularity conditions and with infinite rounds of communication, consensus eventually emerges. However, DeGroot (1974)'s approach implies the repetition of information and, thus, a sub-optimal aggregation of the information available to the network. To update their estimates correctly, rational agents should recount the sources of all information that contributed to forming their beliefs, the beliefs of those to whom the agent listens and

[1] Even if our analysis is mainly confined to butterfly strategies, other option strategies could be considered to reflect a different degree of precision required in the estimation of the unknown parameter for an agent to make the right decision (see Section 2).

[2] The listening set of an agent comprises the individuals within the network to whom the agent listens. To keep the problem tractable, we exclude the possibility that new information could be received before a decision has been made.

[3] This phenomenon refers to the circumstance that, over time, individuals' opinions on a multidimensional set of issues can be well approximated by a simple linear (i.e. unidimensional) structure, where an individual's position on the line determines the individual's position on all issues.

those to whom they listen, and so on. Thus, it would be extremely difficult to implement a rational updating in a large network, where beliefs are derived from many (overlapping) sources over an extended period. DeMarzo et al. (2003) introduces a concept of bounded rationality, named *persuasion bias*, to justify the repetition of information. In every communication round, each agent acts as though the beliefs spread by all agents were always derived from new and independent observations. Experimental investigations of *persuasion bias* within social networks can be found, among others, in Corazzini et al. (2012) and Brandts et al. (2015) under mild conditions. While the results in Brandts et al. (2015) are consistent with the theoretical framework, Corazzini et al. (2012) find that the most influential agents are not those with more outgoing links, as predicted by the *persuasion bias* hypothesis, but those with more incoming links. Other works that exploit *persuasion bias* for the transmission of political information are Enikolopov et al. (2011), Barone et al. (2015), and Martin and Yurukoglu (2017), while Hong et al. (2004) analyzes *persuasion bias* in the stock market. Golub and Jackson (2010) consider a variation of DeMarzo et al. (2003)'s model and highlight the pervasiveness of *persuasion bias* by excluding the fact that a simple and boundedly rational updating rule leads to an accurate estimate of the unknown parameter, even if network grows large, unless no agent receives disproportionately high attention. Our model, whilst adopting a non-Bayesian opinion aggregation, departs from this literature in considering rational agents who dynamically update the weights assigned to their peers' beliefs given the information made available at every communication round. The model retains the spirit characterizing the DeGroot model because agents revise their weights and beliefs by evaluating the weights assigned by their peers at every communication round. However, the weights spread throughout every communication round now represent an actually new piece of information, that is, the evaluation of a different ability of each agent. Thus, repetition of information and *persuasion bias* are ruled out at the root. Unlike our view of the informative content of the network, the idea of time-dependent weights has been widely proposed in the literature. DeMarzo et al. (2003) assume that the weights each agent assigns to her own peers as a whole might change proportionally over time as a consequence of her increasing self-confidence. More interestingly, Rapanos (2023) proposes a model of dynamically updated weights, where the weights assigned to the agents vary depending on the sources of information they contact over time. Agents with an initially low weight could be assigned a higher weight once they gather new information from a highly reliable source of information. Conversely, in our model, a low weight initially assigned to an agent might be revised after a new communication round not because of interaction with more valuable sources of information but because a different ability is considered. An agent who is not reliable when evaluating the unknown parameter could be trustworthy when the reliability of the initial estimates is considered. Polanski and Vega-Redondo (2023) analyzes the co-evolution of networks and opinions and suggests homophily as a descriptive postulate leading, under some topological conditions, to polarization. The weight of each link within a network should match the similarity of opinions of the connected agents. The role of homophily in leading to polarized opinions has been widely analyzed in many other papers, such as Melguizo (2019), and can be traced back to the seminal work by Lazarsfeld and Merton (1954). In our model, homophily could arise if agents do not expect to receive much future information and are therefore willing to give a positive weight only to peers with close estimates. However, we depart from this literature because our proposed model is not shaped by the observation of any social pattern or communication bias, and the final outcome of agents' interactions depends only on the parameters that define each agent's attitude toward both future information and estimating errors. As a consequence, both consensus and polarization are allowed as final outcome of any social interaction. The paper is organized as follows: in Section 2, we present the main features of our model and, in particular, our new interpretation of the informative content of the network. Two illustrative examples are also provided. In section 3 the agents' belief dynamics is described. In Section 4.1, a sufficient and a necessary condition for convergence are outlined. In addition, unidimensionality of opinions and a special case with weights remaining constant over time are briefly discussed. Finally, Section 5 is devoted to some empirical examples that demonstrate how the model is able to properly mimic real situations.

2. The model

Consider a finite set of agents $\mathcal{N} = \{1, 2, \dots, N\}$, indexed by i, j or h, which constitute a social network. The initial network is described as a directed graph indicating whether agent i "listens to" agent j. This graph, otherwise named a listening structure, is assumed to be exogenous and may correspond, for example, to geographical proximity or social or hierarchical relationships. An agent i's listening set is defined as $\mathcal{S}_i \subseteq \mathcal{N}$, and for every agent j, we define an indicator parameter $q_{ij} \in \{0, 1\}$ that specifies whether agent j belongs to $\mathcal{S}_i$ or not, with $q_{ij} = 1$ if and only if $j \in \mathcal{S}_i$. We assume that each agent listens to herself, that is $i \in \mathcal{S}_i$ for every $i \in \mathcal{N}$. Let $\theta \in \mathbb{R}^L$ be a L-dimensional unknown parameter (or state of nature) that each agent has to estimate to make a decision. For the sake of simplicity, two assumptions are introduced: first, we restrict our analysis to a fully connected network with $\mathcal{S}_i = \mathcal{N}$ for every agent $i \in \mathcal{N}$; second, we consider a unidimensional unknown parameter $\theta \in \mathbb{R}$.

2.1. Agents' utility function

The utility of the agents is contingent upon the accuracy of their decisions or, which is equivalent, on the degree of precision of their estimates of the unknown parameter θ. Thus, their utility is modeled as inversely proportional to the absolute estimation error, which is defined as the absolute difference between the agent's estimate x_i^0 of the unknown parameter and its actual value θ. Furthermore, the utility function is bounded from below: if the estimation error exceeds a given threshold δ, the resulting utility attains its minimum value. Agent i's utility function $u_i : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ is then defined as follows:

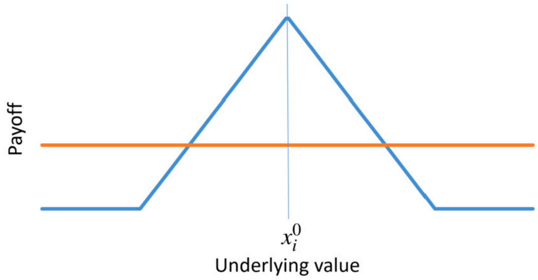

Fig. 1. Payoff of a butterfly spread strategy.

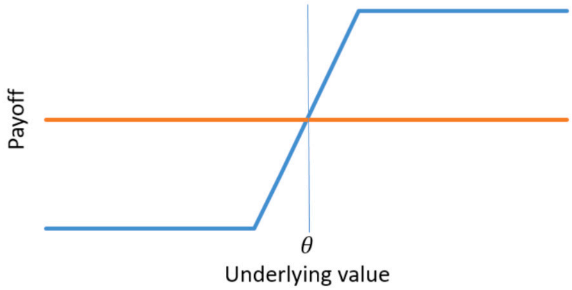

Fig. 2. Payoff of a bull spread strategy.

$$u_i\left(x_i^0,\theta\right)=\begin{cases}0 & \text{if} \quad |x_i^0-\theta|\geq\delta\\ \theta-a & \text{if} \quad 0\leq x_i^0-\theta<\delta\\ c-\theta & \text{if} \quad 0<\theta-x_i^0<\delta\end{cases} \tag{1}$$

where a and c are real positive parameters with $a<x_i^0<c$ and $x_i^0-a=c-x_i^0=\delta$. An agent's utility is thus analogous to the payoff of a financial derivative instrument designated as *butterfly spread* [4] (see Fig. 1) when its middle strike is equal to x_i^0. Therefore, when a network is considered, the adoption by agent i of a parameter estimate x_j^0 shared by some agent j is equivalent to purchasing a butterfly option with agent i's estimate $x_{i,t}^0$ as underlying asset and middle strike equal to x_j^0. Agent i's payoff will be maximal only if x_j^0 coincides with the eventually revealed true value of θ, while the value of the parameter δ can be properly calibrated to reflect the agent's payoff sensitivity with respect to an estimate error. It is well known (see, for example, Blyth (2014)) that, given a value of δ arbitrarily close to zero, the payoff of the butterfly approximates the (risk neutral) probability of the underlying being between the strikes of the two long call options composing the butterfly. Therefore, given suitable values of δ, agents do trade probabilities when trading butterflies, in particular the risk neutral probabilities of profitable values of the underlying random variable.

2.2. Agents' utility: a generalization of the approach

Alternative functional forms to equation (1), which similarly capture the payoff structures of option strategies like the butterfly spread, could also be considered. More specifically, the selection of a specific functional form for agent i's utility is contingent upon the degree of estimation precision required to make an accurate decision. An agent's utility mimics the payoff structure of a ***bull (bear) spread strategy***[5] (see Fig. 2) when an accurate decision is made, provided that the estimated parameter exceeds a lower (or upper) bound irrespective of the degree of precision. The payoff then decays linearly as the value of x_i^0 deviates from this bound, ultimately

[4] A butterfly spread is a neutral-outlook option strategy whose payoff depends primarily on the underlying financial asset's price S_T at the expiration date T relative to the strike prices of the options involved. A butterfly spread, using call options, consists of buying one call option with a lower strike price $a=\theta-\delta$, selling two call options with a middle strike price θ, and buying one call option with a higher strike price $c=\theta+\delta$ (see Hull (2006)). All options have zero cost and an expiration date of T. The butterfly's final payoff is defined as follows:

$$\mathit{Payoff}=\max\left(S_T-a,0\right)-2\cdot\max\left(S_T-\theta,0\right)+\max\left(S_T-c,0\right) \tag{2}$$

Thus, the butterfly spread is most profitable when the underlying asset's price is close to the middle strike price at expiration, and it has limited risk on both the upside and downside. Furthermore, a narrow spread (i.e., strike prices that are closer together) tends to create a steeper peak. This means the maximum profit is reached quickly, but the range of values at which the strategy is profitable is smaller. Conversely, a wider spread (i.e., strike prices that are further apart) results in a flatter peak and a gentler slope. This allows for a broader range of values where the strategy is profitable, but the maximum profit is lower. Note that equations (1) and (2) coincide given $S_T=\theta$.

[5] See Hull (2006) for a formal definition.

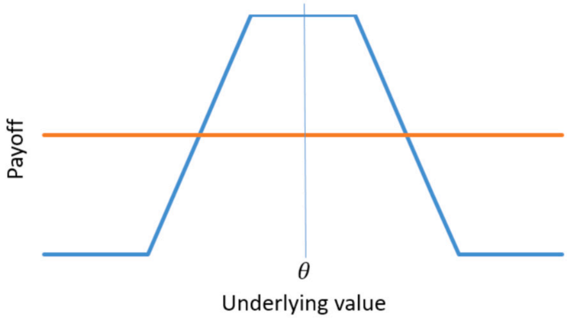

Fig. 3. Payoff of a condor strategy.

reaching $l \leq 0$ when the estimation error exceeds a threshold of δ. Similarly, a *condor strategy*[6] properly describes an agent's payoff when moderate errors in estimation are permissible without significantly impacting the overall outcome (see Fig. 3).

2.3. Future evolution of agents' initial estimates

In financial markets, a butterfly spread is constructed on a financial asset whose value evolves over time according to a specified stochastic process. At first glance, the estimated value x_i^0 of the unknown parameter θ, representing the underlying in our setting, appears to diverge considerably from the price of a financial asset. However, as a stock's price evolves over time according to the information received by the market, the estimated vector of unknown parameters or weights changes over time as agent i receives new information.

Assumption 2.1. Once consensus has been reached within the network, or agent i has made her decision, her estimate of the unknown parameter $x_{i,t}^0$ evolves over time according to a discrete compound Poisson process[7] with intensity λ_i^0 and, eventually, converges to the actual value of the parameter θ.

The compound Poisson process enables the modeling of the future flow of relevant information that each agent i might receive about the true value of the parameter, provided that both the relevance of each new piece of information and its time of occurrence can be modeled as random variables. Note that the value of the parameter λ_i^0 is agent-specific, reflecting the expected amount of future information. Consequently, its value should be inversely proportional to the agent's degree of confidence in their estimate. Intuitively, an accurate estimate of λ_i^0 would require knowledge of both the structure of the network of which agent i will be part of in the future (because we allow for new sources of information) and the reliability of the initial estimate x_i^0. If the quality of the initial estimate is poor, the probability that the agent will receive relevant new information in the future will be high. Furthermore, a natural question arises: which weight should be assigned to future information? It must be noted that *a priori*, an agent i whose reliability is itself a random variable could not know even her future sources of information. Therefore, it seems that no weight can be assigned to future information in advance. However, a compound Poisson process also allows us to take reliability into account because the dimension of the jump is itself a random variable, and the random value of any jump can reflect both the impact of the received information and its quality. Furthermore, we assume that there is a future point in time when agents learn the true value of the parameter because all the information agents gather over time, or receive as a consequence of their choices, eventually reveals or at least refines the initial estimate of the unknown parameter. This assumption, which is crucial to price an option strategy, seems reasonable. As an example, consider a college student's choice of academic curriculum. The value of the unknown (multidimensional) parameter might represent the weights, in percentage terms, of all subjects within the curriculum. The closer the estimate of the parameter to its optimal value, the higher will be the student's utility. Initially, the student makes a choice given the information gathered within her network. However, upon choosing the curriculum, the student could receive new information, such as unexpected exam results or feedback from new professors or other students that enlarge the initial network. Then, as her academic career proceeds, the flow of information eventually unveils the true value of the unknown vector. Similarly, in a political election, one could consider a bi-dimensional unknown parameter, with the first dimension representing the matching between the political program and the voter's preferences, and the second dimension the degree of reliability of the nominee. Then, voters decide which candidate should get their preference, given the information available about the candidates and derived from (and shared with) different sources, such as family members, friends, and social networks. However, after the election, voters could change their minds, given how the nominee actually acts and fulfills her electoral pledges. By the end of the nominee's time in office, the true value of the unknown parameter will be actually known.

[6] See Hull (2006) for a formal definition.

[7] A compound Poisson process represents an extension of a Poisson process, designed to model situations in which not only do events occur randomly over time, but each event also has a random "size" or "impact." The events occur independently, and the number of events in any given time period follows a Poisson distribution, with events happening at a constant average rate λ_i^0 per unit of time. The total outcome of the process is the cumulative sum of the sizes of all the events that have occurred, see Appendix A.

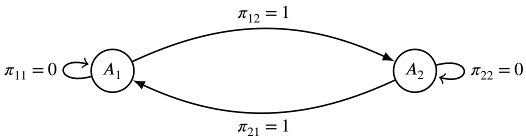

Fig. 4. Two agent network.

2.4. Examples

Two examples are here proposed to justify, in our setting, the assumption of self-confident agents (i.e. $\pi_{ii}^0 \neq 0$ for every $i \in \mathcal{N}$) and to show how multiple rounds of communication are a natural part of social interaction without necessarily implying repetition of information. First, consider a simplified network with only two agents indexed by i with $i = 1, 2$ (see Fig. 4) who, after some communication round $\hat{t}$, must choose one of two mutually exclusive alternatives depending on the estimated value $x_i^{\hat{t}}$ of an unknown parameter $\theta \in \mathbb{R}$. We assume that agents' initial estimates are, respectively, $x_1^0 = 1$ and $x_2^0 = 0$ and that the *listening matrix* T is constant over time:

$$T = \begin{pmatrix} \pi_{11}^0 & \pi_{12}^0 \\ \pi_{21}^0 & \pi_{22}^0 \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$$

Given T, agents' beliefs would move from one (zero) to zero (one) back and forth through communication rounds. Thus, dropping the assumption of self-confidence excludes consensus[8] and leads, in our setting, to a striking contradiction: by construction, each agent assigns a zero weight to her initial estimate. However, after the second communication round, that is, upon knowing that agent j considers x_i^0 trustworthy, agent i assigns weight $\pi_{ij}^1 = 1$ to agent j. In other words, according to agent i, agent j correctly considers x_i^0 a reliable estimate of θ.[9] This contradicts her initial evaluation of the quality of x_i^0. Conversely, one might have expected that, after the second communication round, upon discovering that they assigned weight one to each other, both agents would doubt their ability to assess, respectively, π_{ij}^0 and π_{ji}^0. If agent i gives weight zero to her own information, then there is no reason for agent j to consider that information as trustworthy, and vice-versa. The weights in π_i^0 and π_j^0 might suggest that agent i is either underestimating the quality of x_i^0 or, symmetrically, overestimating the informative content of x_j^0. Therefore, the weight π_{ii}^1 should be positive for any agent i. This example, while representing a degenerate case, also seems useful in justifying multiple communication rounds: after the first communication round, an agent who is certain about her own initial weights should not enter any other communication round. An exception would be the case where the agent's opponent does serve as a reliable source of information. Further interaction among the agents can be justified only if agents were not confident about their assigned weights and believe other agents have valuable information about them. Now, consider the network represented in Fig. 5 where agent J represents a journal with a clear political slant[10] and agents $A_i = \{1, 2\}$ the readers. The initial weight assigned to the journal's article by the agents reflecting on the journal's *ability* to evaluate a specific issue is $\frac{3}{8}$. Zero weight is implicitly placed by J on both agents, because agent J does not even listen to them.[11] The initial *listening matrix* T is then defined as:

$$T = \begin{pmatrix} 1 & 0 & 0 \\ \frac{3}{8} & \frac{1}{4} & \frac{3}{8} \\ \frac{3}{8} & \frac{3}{8} & \frac{1}{4} \end{pmatrix} \tag{3}$$

It is easy to observe that if the newspaper confirmed a weight 1 to itself over time, and the other two agents kept assigning some strictly positive weight to J at every communication round, then the agents' beliefs would eventually converge to the journal's belief x_J^0. According to *persuasion bias* this happens because agents A_1 and A_2 do not realize that agent J is repeating the same information over time. Note that in this case the value of the unknown parameter spread by agent J is constant over time and, thus, repetition of information should be obvious and interaction should stop. Conversely, in our setting, the information provided by J at the second communication round is actually new. Agent J reveals that its evaluation of the parameter is, up to its knowledge, correct. At the third communication round, the journal confirms that the piece of information communicated at the second communication round is also correct and so on. The journal finally emerges as a *guru* within the network. This mimics the persuasion phenomenon, with a

[8] Given a Markov chain with constant transition matrix T, dropping the hypothesis $\pi_{ii}^0 > 0$ for any $i \in \mathcal{N}$ is equivalent to dropping the property of aperiodicity of the states which is a necessary condition for convergence. Reconsidering the example in "Markovian terms", agents' initial estimates correspond to the two states of the Markov chain, with transition probabilities given by the assigned weights. Both states have periodicity two.

[9] Recall that π_{ij}^1 reflects agent i's assessment of agent j's ability to evaluate the proficiency at estimating the unknown parameter.

[10] This example has already been proposed in DeMarzo et al. (2003).

[11] An agent could assign zero weight to a source of information due to a lack of connection or confidence. However, in this case, zero weights are not misleading because, although they stem from a lack of connection, we assume agent J is actually self-confident and would assign, in any case, a zero weight to any other agent. Note that in this case we depart from the general assumption of a fully connected network. However, the example would still easily work if we considered a fully connected network and if zero weights were replaced by strictly positive weights arbitrarily close to zero.

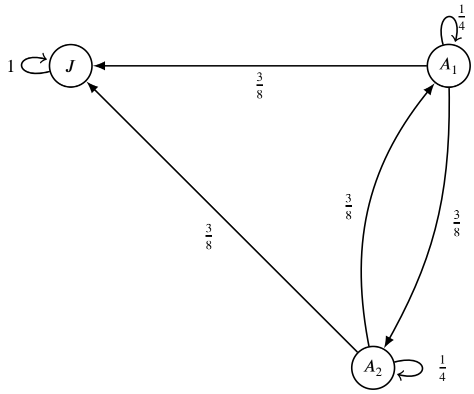

Fig. 5. Example 2.

slightly different perspective because we are focusing on the reliability and the reputation of the sources that emerge over time. It is clear that the higher the journal's overall reputation (i.e., the higher the weights assigned to it over time), the faster the speed of convergence.

3. Beliefs dynamics

3.1. First communication round

3.1.1. Estimate of agents' precision

Formally, communications and estimate updates in our model are as follows. Each agent i possesses some possibly different initial information x_i^0 regarding the value of the unknown parameter θ. In the first communication round, agents truthfully report their initial evaluations x_i^0 and listen to other agents' estimates x_j^0 with $j = 1, \dots, N$. Assigning a positive weight $\pi_{ij}^0 \in [0, 1]$ to the available estimate x_j^0 corresponds to agent i *allocating a share* π_{ij}^0 *of the available wealth* on the butterfly B_{ij}^0 with a center strike x_j^0 and width δ_i. If agent i considered agent j as a completely trustworthy source of information she would give weight $\pi_{ij}^0 = 1$ to the estimated value x_j^0. This corresponds to investing all the wealth in the butterfly with the center strike x_j^0. Conversely, if an agent does not regard any source of information as fully reliable, it would be reasonable to assign some weight to a collection of estimates, thus to a set of butterflies. We suggest that each agent defines the weights by constructing a (virtual) portfolio of butterflies as the solution to a mean-variance portfolio optimization problem[12]:

Proposition 1. *Given the set $\mathcal{B}_i^0$ of available butterflies, the precision level π_{ij}^0 assigned by agent i to each agent j at the first communication round is represented by the weight assigned to butterfly B_{ij}^0 in $\mathcal{B}_i^0$ determined as the outcome of a mean-variance optimization problem.*

The choice of the mean-variance optimization seems natural because we are considering a risk-averse decision maker who aims to maximize the portfolio's expected return and to minimize its variance. Unlike the former goal, the latter requires further discussion. Modern portfolio theory or mean-variance analysis, as introduced by Markowitz (1952), considers risk-averse investors who prefer outcomes with low uncertainty to those with high volatility even if the expected outcome of the latter is equal to or greater than the more certain outcome. Accordingly, we consider a risk-averse decision maker who prefers to avoid making choices with highly uncertain outcomes. Risk-averse investors prefer to diversify among financial instruments. Similarly, risk-averse agents who are uncertain about their estimates, and who, therefore, are willing to account for future information, prefer to diversify among different opinions by assigning positive weights to several agents.

3.1.2. The updating process

Each agent i, upon learning the estimates of the other agents in the listening set $\mathcal{S}_i$, updates the estimate of the parameter by computing a weighted average with weights π_{ij}^0 which are the solution of the portfolio optimization described in the previous paragraph. The updated estimate, designated as x_i^1, is defined as follows:

$$x_i^1 = \sum_j q_{ij} \pi_{ij}^0 x_j^0 \tag{4}$$

The updating rule (4) can be expressed more concisely in vector notation. Let x^0 represent the vector of estimates with elements x_j^0 and let T^0 denote the listening matrix with elements $T_{ij}^0 = q_{ij} \pi_{ij}^0$. Then equation (4) can be written as follows:

[12] The analytical description of the portfolio optimization process is reported in Appendix A.

$$x^1 = T^0 \cdot x^0 \tag{5}$$

3.2. Second communication round

In the second communication round, each agent truthfully discloses her evaluations of the reliability of the other agents, as represented by the vector of weights π_{ij}^0 assigned to the butterflies that comprise the optimal portfolio defined in the initial communication round. In accordance with the outlined theoretical framework, we assume that each element π_{ij}^0 of the vector π_i^0 evolves over time according to a compound Poisson process[13] reflecting the information agent i might receive about the reliability of her information sources:

Assumption 3.1. Once consensus has been reached within the network, or agent i has made her decision, her estimate of any other agent reliability $\pi_{ij,t}^0$ evolves over time according to a discrete compound Poisson process with intensity λ_{ij}^1 and, eventually, converges to the actual value of the vector-valued parameter π_i^0.

Let $\mathcal{B}_i^1$ be the new set of butterflies in which agent i could invest after this second interaction. For each butterfly B_{ij}^1 in $\mathcal{B}_i^1$, with $j = 1, \dots, N$, the underlying asset is now represented by the vector-valued random variable $\pi_{i,t}^0$ while the middle strike is set at π_j^0. Once again, each agent i assigns a weight π_{ij}^1 to each butterfly B_{ij}^1 in $\mathcal{B}_i^1$ by composing a (virtual) optimal portfolio obtained as a solution to a mean-variance optimization problem. Agents' beliefs, after the second communication round, are then represented as follows:

$$x^2 = T^1 \cdot T^0 \cdot x^0 \tag{6}$$

given the matrix T^1 with elements $T_{ij}^1 = q_{ij}\pi_{ij}^1$

3.3. Subsequent communication rounds

The vector of weights π_i^1 will be shared by agent i at the third communication round and will become the underlying of a new set of butterflies $\mathcal{B}_i^2$, each centered on a different vector π_j^1 and with the underlying of a vector $\pi_{i,t}^1$ of random variables $\pi_{ij,t}^1$ each indexed by j and driven by a compound Poisson process with intensity parameter λ_{ij}^2. A new (virtual) optimal mean-variance portfolio including just butterflies in $\mathcal{B}_i^2$ is then composed by each agent and agents' beliefs are updated accordingly. The same reasoning applies to any subsequent communication round. Therefore, we end up with countably many (virtual) portfolios of butterflies for every agent, with each portfolio constructed after every communication round. Note that, in principle, there is no reason to impose any correlation between the Poisson processes driving the weights shared at different communication rounds because we are referring to different abilities. However, one could assume that the value of the intensity parameter λ_{ij}^t with $t = 1, 2 \dots$ decreases or at least does not increase with time because, as the number of communication rounds increases, agent i may be less likely to receive information about the specific ability to be evaluated. In general terms, we could write the following:

$$\lambda_{ij}^t = \lambda_{ij}^1 \, k_i(t) \quad \text{for } t > 1 \tag{7}$$

where $k_i(t)$ is some non-increasing and non-negative function of t for every agent i. However, if the value of λ_{ij}^t decreased too rapidly, this would result in increasingly confident agents and, consequently, put consensus at risk.

4. Opinion convergence

In our context, a crucial issue is to determine the conditions under which, given the sequence of stochastic matrices resulting from each communication round, the sequence of their left products converges to a rank-one matrix.[14] We will refer to a set of matrices $\mathcal{P}$ satisfying this property as a consensus set:

Definition 4.1. A set $\mathcal{P}$ of stochastic matrices is a consensus set if, for every sequence of matrices $P_1, P_2, P_3 \dots$ whose elements belong to $\mathcal{P}$ and for every initial state x_0, the sequence of states defined by $x_t = P_t \cdot P_{t-1} \dots P_1 \cdot x_0$ converges to a vector whose entries are all identical.

[13] After each jump, the resulting vector of non-negative weights must always have the sum of its components equal to one. In other terms, we assume that π_i^0 could just move within a $N-1$ dimensional simplex and, therefore, that the weights may need to be normalized. Alternatively, we could assume that the relative weights that each agent i assigns to other agents providing information follow a compound Poisson process over time. Still, we must normalize the weights to ensure they are non-negative and sum up to one.

[14] Necessary and sufficient conditions for an infinite left product of stochastic matrices converging to a rank-one matrix have been widely studied in the literature; see, e.g., Tsitsiklis et al. (1986), Olfati-Saber and Murray (2004), Ren and Beard (2005), Cao et al. (2008), Olshevsky and Tsitsiklis (2009), Cao et al. (2008), and Egerstedt et al. (2012).

The literature on the characterization of a consensus set can be traced back, at least, to the work of Wolfowitz (1963) in which the class of stochastic, indecomposable, aperiodic (SIA) matrices was first introduced:

Definition 4.2. A matrix P is called SIA, i.e., stochastic, indecomposable and aperiodic, if it is stochastic and the limit

$$Q = \lim_{n\to\infty} P_n \tag{8}$$

exists and all the rows of Q are the same.

A classical result in Wolfowitz (1963) states that any product of transition matrices from $\mathcal{P}$ converges to a rank-one matrix if and only if every product of matrices in $\mathcal{P}$ is SIA. Informally, the chain composed by such matrices in $\mathcal{P}$ forgets its distant past. This theorem has found many applications (see, e.g., Ren and Beard (2005) and Sarymsakov (1961)) and guarantees convergence of the modeled stochastic system in every possible scenario or switching between matrices in $\mathcal{P}$. More generally, we will make use of the following result:

Theorem 4.1. *Let $\mathcal{P}$ be a compact set of $n \times n$ stochastic matrices. The following conditions are equivalent.*

1. *$\mathcal{P}$ is a consensus set.*
2. *For each integer $k \geq 1$ and any $P_i \in \mathcal{P}$ with $1 \leq i \leq k$, the matrix $P_k \cdot P_{k-1} \dots P_1$ is SIA.*
3. *There is an integer $\upsilon \geq 1$ such that for each $k \geq \upsilon$ and any $P_i \in \mathcal{P}$ with $1 \leq i \leq k$, the matrix $P_k \cdot P_{k-1} \dots P_1$ is scrambling.*
4. *There is an integer $\mu \geq 1$ such that for each $k \geq \mu$ and any $P_i \in \mathcal{P}$ with $1 \leq i \leq k$, the matrix $P_k \cdot P_{k-1} \dots P_1$ has a column with only positive elements.*
5. *There is an integer $\alpha \geq 1$ such that for each $k \geq \alpha$ and any $P_i \in \mathcal{P}$ with $1 \leq i \leq k$, the matrix $P_k \cdot P_{k-1} \dots P_1$ belongs to the Sarymsakov class.*

While condition (2) provides a necessary condition for convergence, imposing the requirement that every matrix in $\mathcal{P}$ must be SIA, conditions (3), (4), and (5) establish that $\mathcal{P}$ must be a set of matrices with at least one column of strictly positive elements or must consist of scrambling or Sarymsakov matrices, which constitutes a sufficient but not necessary condition for convergence.

4.1. Positive matrices

We first consider a proper subset of the SIA class of stochastic matrices, namely positive stochastic matrices. Although this represents a very special case of a network in which all agents assign positive weights to each other, it allows for the formulation of a sufficient condition for opinion convergence with an interesting financial interpretation.

Definition 4.3. A stochastic matrix A is positive if and only if $a_{ij} > 0$ for all i, j

Proposition 2. *Let $A_1, A_2, A_3 \dots$ be an infinite sequence of matrices where each matrix A_i is a positive stochastic matrix. Then, the left product $A_n A_{n-1} \dots A_1$ converges to a rank-one matrix as $n \to \infty$.*

In other words, consensus is reached if every agent within the network is regarded as a trustworthy information source or, equivalently, if at every communication round all available option strategies receive a positive weight in each agent's portfolio. Defining necessary and sufficient conditions that ensure non-negative (positive) portfolio weights is a well-known problem in portfolio theory. We refer to this literature and, in particular, to Best and Grauer (1992) who derive an easily computable necessary and sufficient condition for a minimum variance portfolio to achieve all non-negative (positive) weights.[15] In order to formalize their argument, consider the optimization problem each agent i solves at the first communication round. Given the collection $\mathcal{B}_i^0$ of N butterflies, let Σ be its variance-covariance matrix, μ its N-dimensional vector of expected returns and e the N-vector whose components are all units. The matrix Σ is positive definite and the value of the parameter γ_i represents the agent's risk tolerance parameter: the higher the value of γ_i, the more tolerant the investor will be to risk. The mean-variance (hereafter MV) problem is:

$$\max_{\pi_i^0} \left\{ \gamma_i \mu^\intercal \pi_i^0 - \frac{1}{2} \left(\pi_i^0\right)^\intercal \Sigma \pi_i^0 | e^\intercal \pi_i^0 = 1 \right\} \tag{9}$$

where π_i^0 is the N-vector of portfolio weights, and $e^\intercal \pi_i^0 = 1$ is the budget constraint. Equation (B.7) in Appendix B implies that every component of the optimal vector of weights $\pi_i^0\left(\gamma_i\right)$ is a linear function of γ_i that is:

$$\pi_{ij}^0\left(\gamma_i\right) = \alpha_{0j} + \gamma_i \alpha_{1j} \quad \text{with} \quad j = 1, \dots, N$$

[15] Green and Hollifield (1992) extended the results in Best and Grauer (1992) to the case where the weights of the minimum variance portfolios lie between upper and lower bounds.

If $\alpha_{1j} > 0$ $(\alpha_{1j} < 0)$, then $\pi^0_{ij}(\gamma_i)$ is increasing (decreasing) in γ_i and will be non-negative, provided that $\gamma_i \geq -\frac{\alpha_{0j}}{\alpha_{1j}}$ $\left(\gamma_i \leq -\frac{\alpha_{0j}}{\alpha_{1j}}\right)$. Let:

$$\gamma^l_i = \max\left\{-\frac{\alpha_{0j}}{\alpha_{1j}} \mid \text{all } j \text{ with } \alpha_{1j} > 0\right\} \tag{10}$$

$$\gamma^u_i = \min\left\{-\frac{\alpha_{0j}}{\alpha_{1j}} \mid \text{all } j \text{ with } \alpha_{1j} < 0\right\} \tag{11}$$

It follows that:

$$\pi^0_{ij}(\gamma_i) \geq 0 \text{ for all } j \text{ with } \alpha_{1j} > 0 \text{ and for all } \gamma_i \geq \gamma^l_i \tag{12}$$

$$\pi^0_{ij}(\gamma_i) \geq 0 \text{ for all } j \text{ with } \alpha_{1j} < 0 \text{ and for all } \gamma_i \leq \gamma^u_i \tag{13}$$

If α_{1j} were equal to zero for one or more assets j, in order to have $\pi^0_{ij}(\gamma_i) \geq 0$, it would be necessary that:

$$\alpha_{0j} \geq 0 \text{ for all } j \text{ such that } \alpha_{1j} = 0 \tag{14}$$

Therefore, given Theorem 1 in Best and Grauer (1992), we can state a necessary and sufficient condition for the transition matrix T^0 to be positive:

Theorem 4.2. *The transition matrix T^0 is positive if and only if for every agent $i \in \mathcal{N}$, given the collection $\mathcal{B}^0_i$ of N butterflies with positive definite variance-covariance matrix Σ, we have $\gamma^l_i < \gamma^u_i$ and $\alpha_{0j} > 0$ for all j such that $\alpha_{1j} = 0$. The entries of matrix T^0 are given by $T^0_{ij} = \pi^0_{ij}(\gamma_i) = \alpha_0 + \gamma_i \alpha_1$ for all γ_i satisfying $\gamma^l_i < \gamma_i < \gamma^u_i$.*

Intuitively, this condition explicitly requires a certain degree of risk-aversion for each agent (see Appendix B) and can be easily extended to every communication round to get a sequence of positive matrices and thus convergence. Interestingly, Best and Grauer (1992) have also proven that positively weighted minimum-variance portfolios may not exist when the number of assets increases significantly. This result can be easily translated to our setting: as the number of agents within the network increases, it becomes more difficult to reach a consensus. As a second condition a positive definite variance-covariance matrix Σ is required. This would exclude arbitrarily small values of the intensity parameter λ^0_i that characterizes the compound Poisson process. When the underlying asset experiences infrequent jumps, its price remains near its initial value x^0_i for a significant period. Consequently, butterflies with central strikes far from x^0_i would not be profitable, leading to zero variance in their payoffs. This intuitively implies that agents should be confident about the future flow of relevant information.

4.2. Sarymsakov matrices

It is important to note, however, that while positive transition matrices are a sufficient condition for convergence, they are not a necessary one. There exist classes of transition matrices that, despite not being positive, still ensure convergence. In the literature, the set of stochastic Sarymsakov matrices, first introduced by Sarymsakov (1961), is the largest known subset of the class of stochastic matrices whose compact subsets are all consensus sets; in particular, the set is closed under matrix multiplication, and the left product of the elements from its compact subset converges to a rank-one matrix.

Definition 4.4. Given a non negative N-dimensional matrix T, for any set $S \subseteq \{1, \ldots N\}$, the consequent function F_T is defined as follows:

$$F_T(S) = \left\{j : \exists i \in S \text{ such that } T_{ij} > 0\right\} \tag{15}$$

Then, a stochastic matrix T is called a Sarymsakov matrix if and only if, for any two disjoint non-empty subsets S and $\tilde{S}$, given $F_T(S)$ and $F_T(\tilde{S})$, either $F_T(S) \cap F_T(\tilde{S}) \neq \emptyset$ (first condition) or $F_T(S) \cap F_T(\tilde{S}) = \emptyset$ and $|F_T(S) \cup F_T(\tilde{S})| > |S \cup \tilde{S}|$ (second condition), where $|S|$ denotes the cardinality of S.

We say that T is a scrambling matrix if for any pair of distinct indices $i, j \in N$, there holds $F_T(i) \cap F_T(j) \neq \emptyset$, which is equivalent to the property that there always exists an index $k \in N$ such that both p_{ik} and p_{jk} are positive. From the preceding definitions, it is clear that a scrambling matrix belongs to the Sarymsakov class. It has been shown that any product of $n-1$ matrices from the Sarymsakov class is a scrambling matrix. Because a scrambling matrix is SIA, any Sarymsakov matrix must be an SIA matrix.

Proposition 3. *Let $A_1, A_2, A_3 \ldots$ be an infinite sequence of Sarymsakov matrices. Then, the left product $A_n A_{n-1} \ldots A_1$ converges to a rank-one matrix as $n \to \infty$.*

Intuitively, $F_T(S)$ is indeed the set of informative sources with influence on the agents in the set S. The definition of a Sarymsakov matrix implies that the sets S and $\tilde{S}$ may have influencing nodes in common or have no influencing nodes in common but that the number of influencers is greater than that of influences. A scrambling matrix is one for which each pair of distinct nodes shares at least one common influencing node. In the context of our analysis, the Sarymsakov matrices emerge when agents demonstrate the willingness to construct diversified portfolios. As previously indicated, this stipulation can be satisfied when a certain level of risk aversion is present. Moreover, the jump processes must possess an intensity parameter that is sufficiently large to generate payoffs from a diverse range of butterflies with varying strike prices over time. Each agent anticipates acquiring new information regarding the unknown parameters, which leads them to assign a positive weight to other butterflies that may prove profitable. If the jumps were too rare or too small, and if most of the portfolio's payoff will be concentrated in butterflies near the current underlying price, there will be poor diversification. In contrast to positive definite matrices, to the best of our knowledge, there is no established theoretical framework that defines the necessary and sufficient conditions to guarantee that the optimization process will result in a transition matrix belonging to the class of Sarymsakov matrices.

4.3. Necessary conditions for convergence

In light of the above discussion, a necessary condition for consensus is straightforward. First, we introduce an assumption about the intensity parameters that characterize the composite Poisson processes in each communication round:

Assumption 4.3. $\sum_{h=1}^{\infty} \lambda_{ij}^{h} \to \infty$

Proposition 4. *If an infinite left product of non-homogeneous transition matrices will converge then:*

1. *Assumption 4.3 is not violated, and*
2. $\max_{i\in\mathbb{N}} \lambda_{ij}^{h}$ *is not arbitrarily close to zero for every* $h = 0, 1, \ldots$.

A violation of one the two conditions in Proposition 4 is a sufficient condition for polarization:

- if 4.1 is violated then λ_{ij}^{h} goes to zero too quickly and the transition matrix will become an identity matrix meaning that the agents become too self-confident at short hand;
- if 4.2 is violated, then agents give a positive weight only to agents with a close opinion, creating opinion clusters when estimates are far apart.

In other words, agents become increasingly self-confident over time, which in turn leads the transition matrix to converge to the identity matrix, leading to polarization. More interestingly, if the values of the intensity parameter of the compound Poisson process are arbitrarily close to zero, this precludes the possibility of receiving future relevant information. Therefore either the agent maintains their opinion assigning a weight of one to themself, or tends to consider only the opinions of others if they are sufficiently close to their own, within a given confidence range. Note that this resembles the setting proposed by Rainer and Krause (2002), where individuals consider the opinions of others only if they are sufficiently close to their own, within a confidence range. This bounded confidence reflects the idea that people tend to disregard opinions that are too far from their own. If the confidence bound is large, agents tend to reach a global consensus: all agents converge to the same opinion. When it is small, agents tend to form clusters: groups of agents converge to different opinions, and multiple stable opinion clusters emerge. The initial distribution of opinions also plays a role in determining the number and size of clusters.

4.4. Uni-dimensional opinions

Unidimensional opinions are a nice implication of De Groot's model as shown in DeMarzo et al. (2003). They refer to the circumstance that, over time, individuals' opinions on a multi-dimensional set of issues can be well approximated by a simple linear (i.e. unidimensional) structure, where an individual's position on the line determines their position on all issues. According to DeMarzo et al. (2003), *"(...) unidimensionality depends on the Markovian structure of updating (whereby old information does not affect updating in a given period save through the formation of prior beliefs entering into that period) and the constant relative weights that an agent gives others over time"*. More interestingly, they conclude that *"(...) there is little reason to believe that other, more general, updating processes would yield long-run linear differences of opinion"*. Conversely, we prove (see Theorem Appendix D.2 in Appendix D) that our model, when consensus is eventually reached, is characterized by long-run unidimensional opinions, even though the relative weights that an agent assigns to others are not generally constant over time. This result suggests that only the Markovian structure of the updating process leads, under mild conditions, to unidimensionality, excluding constancy of relative weights as a necessary requirement.

4.5. Time constant weights

As already pointed out, in our model weights vary over time because a different ability is considered at every interaction. Time-constant weights are a special case in which, at every communication round, each agent has an identical evaluation of her neighbors' reliability. Intuitively, as the number of communication rounds increases, it may become increasingly difficult for an agent to assess the corresponding ability of other agents. Thus, time-constant weights could reflect a kind of bounded rationality. However, we prove (see Appendix C) that they could also arise under the assumption of perfect rationality as a stationary state of the iterative process that describes the interactions among the agents over time. For clarity, let us provide a brief overview of the iterative process: at each communication round $t > 1$, each agent i solves a portfolio optimization problem by assigning a weight to each available butterfly. The solution is a vector of weights π_i^{t-1} in an N-dimensional simplex Δ^N. At communication round $t+1$, each vector of weights π_i^{t-1} with $i = 1, \dots N$ defined after communication round t becomes the center of a new butterfly. A new optimization problem must be solved by each agent i by assigning a weight to each of the new butterflies. The optimization problem faced by each agent is basically the same at every communication round. What changes, from time to time, are the center strikes of each butterfly while any other parameter is assumed to be constant. Therefore, we claim that the solution to each of the $N \times (T-1)$ optimization problems (one for each agent $i = 1, \dots, N$ at each round $t = 2, \dots T$) is the solution of a general parametric optimization problem for a given set of parameters.[16]

5. Empirical case studies

In this section, we propose some examples to show that our model, when applied to practical situations, leads to results consistent with intuition. In particular, the examples highlight the crucial role of the intensity parameters in determining the outcome of the social interaction. We always consider a fully connected network with four agents sharing their opinions on two issues. For the sake of simplicity, we assume that an estimation error affects payoffs identically across issues and agents. The examples differ in terms of the agents' initial estimates and their self-confidence which is allowed to change over time (for details see Appendix E). Both consensus and polarization of opinions might emerge as result of social interaction. Consensus might have different characterizations with agents converging to a middle opinion (see Example 5.1 and Example 5.3) or following a dominant one (see Example 5.4). On the other hand, interaction could result in polarized opinions when agents stick to their estimates or become increasingly self-confident over time (see Example 5.2). Finally, a kind of herd behavior might also emerge when we consider uninformed agents that, having no reliable information, are willing to follow other agents' ideas (see Example 5.5).

5.1. First example: a general case

Let us assume that agents' initial estimates of the unknown parameters $\theta \in \mathbb{R}^2$ show no clusters:

$$x^0 = \begin{bmatrix} 100 & 40 \\ 110 & 100 \\ 150 & 80 \\ 170 & 50 \end{bmatrix}$$

Each estimate $x^0_{i,l}$, for every agent i and every issue l, evolves according to a compound Poisson process with intensity parameters defined by matrix λ^0 with generic element $\lambda^0_{i,l}$:

$$\lambda^0 = \begin{bmatrix} 0.3 & 0.2 \\ 0.4 & 0.4 \\ 0.5 & 0.5 \\ 0.4 & 0.4 \end{bmatrix}$$

At the first communication round, every agent's most profitable butterfly is centered at their own initial estimate (see Fig. 6). In principle, agents would invest mainly in butterflies with a high expected payoff, even if, given risk-averse agents, a positive weight could also be assigned by agent i to butterflies centered at other agents' opinions just to reduce the strategy risk from a portfolio selection perspective. As a result, agents assign positive weights pairwise:

$$T^0 = \begin{bmatrix} 0.73 & 0.27 & 0 & 0 \\ 0.45 & 0.55 & 0 & 0 \\ 0 & 0.54 & 0.46 & 0 \\ 0.49 & 0 & 0 & 0.51 \end{bmatrix} \quad x^1 = \begin{bmatrix} 102.7 & 56.3 \\ 105.5 & 72.9 \\ 128.5 & 90.8 \\ 135.7 & 45.1 \end{bmatrix}$$

After the first interaction, set $\lambda_i^{t-1} = \lambda$ with $\lambda = 0.1$ for every $i \in \mathcal{N}$ and for every $t > 1$. The value of λ implies that agents are more likely to receive information about the true value of the unknown parameter θ rather than about the reliability of all agents. Still, this ensures that every agent i at every communication round considers at least another agent's estimates to be worthy and thus assigns to that agent a positive weight. In other terms, no agent excludes the possibility of receiving, in the future, new information supporting some other agent's opinions. Consider the second communication round: each butterfly β^1_{ij} available to agent i is centered

[16] In other terms, the solution at any time $t > 1$ to each optimization problem is just a function of all the parameters. Because we consider, an objective function that is continuous with respect to the parameters for each agent i, we can apply the Maximum theorem, which provides conditions for a parametric optimization problem to have, as a solution, an upper hemicontinuous correspondence C^* with non-empty and compact values. Thus, the Kakutani's theorem applies and C^* admits a fixed point. Clearly, if a fixed point were chosen by the agents at some point in time, both the underlying of the butterflies and their weights in the optimal portfolio would be fixed from that time on, and the transition matrix would be constant as in DeGroot (1974)'s model, which can be regarded as a special case of our model.

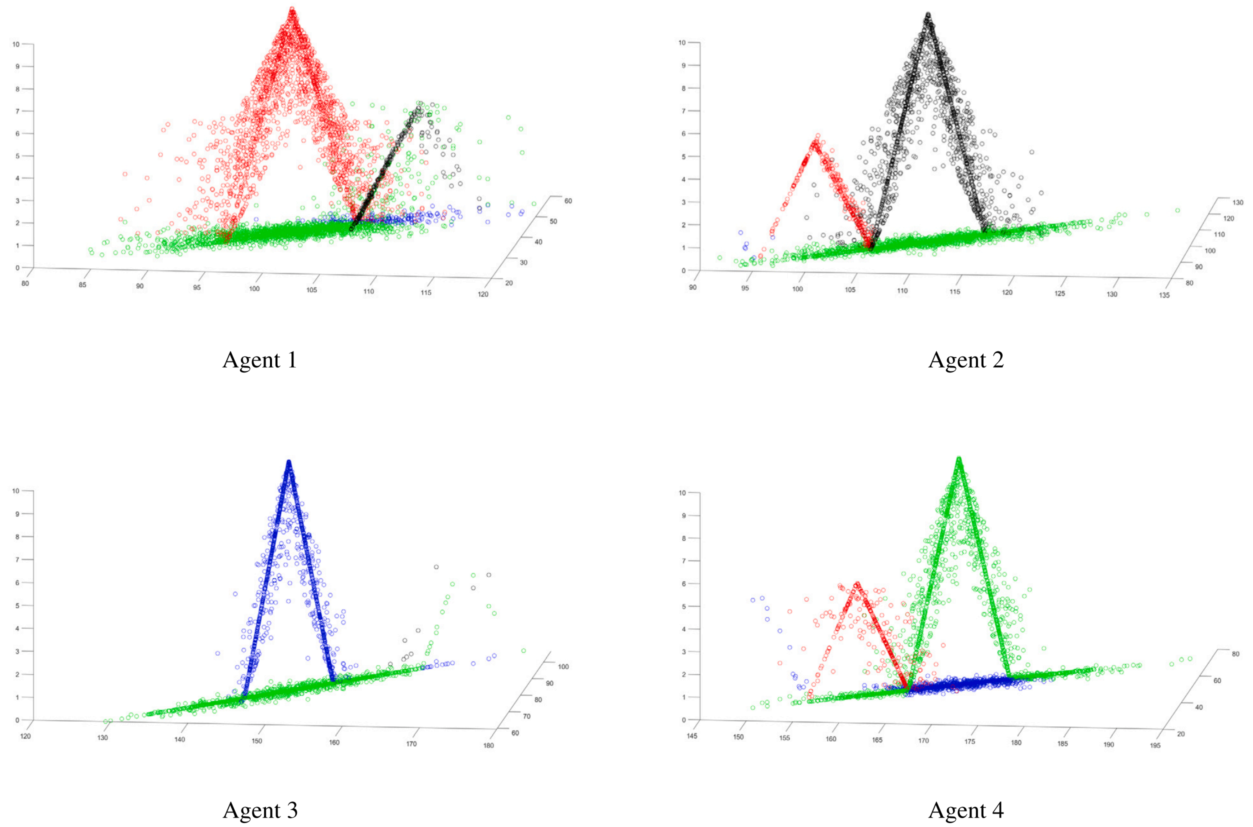

Fig. 6. Butterflies' payoffs after the 1-st communication round. Red color refers to agent 1, black to agent 2, blue to agent 3, and green to agent 4. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

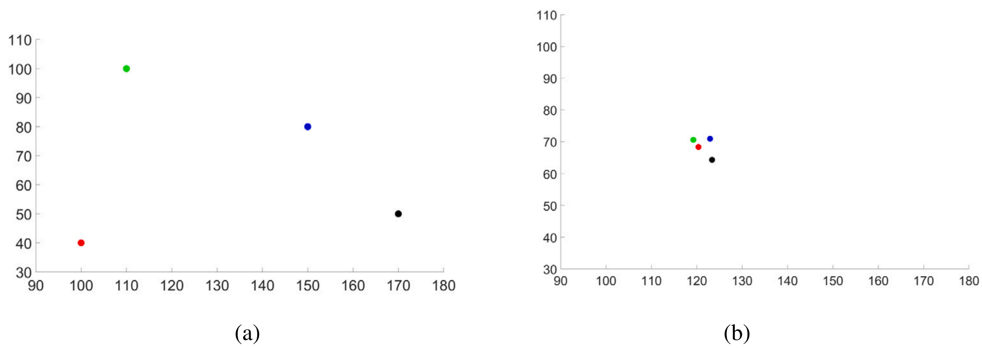

Fig. 7. First example: (a) agents' initial opinions, (b) agents' opinions after the 10-th iteration.

at the vector of weights π_j^0 that agent j computed in the previous step. Then each agent assigns each butterfly a weight by solving a new portfolio optimization problem. The resulting *listening matrix* T^1 and the updated opinions x^2 are as follows:

$$T^1 = \begin{bmatrix} 0.61 & 0.39 & 0 & 0 \\ 0.31 & 0.69 & 0 & 0 \\ 0.01 & 0 & 0.99 & 0 \\ 0.01 & 0 & 0 & 0.99 \end{bmatrix} \quad x^2 = \begin{bmatrix} 103.8 & 62.8 \\ 104.6 & 67.8 \\ 128.1 & 90.5 \\ 135.6 & 45.5 \end{bmatrix}$$

After every communication round, the opinions get closer and eventually converge; see Fig. 7. Note that, not surprisingly, after ten iterations, agents' estimates are arranged almost in a line.

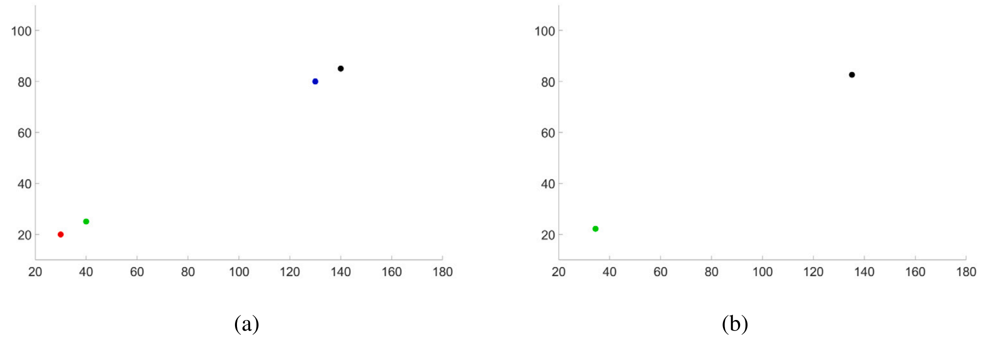

Fig. 8. Second example: (a) agents' initial opinions, (b) agents' opinions after the 10-th iteration.

5.2. Second example: two clusters and increasingly self-confident agents

Consider a network with two clusters, represented by [Agent 1 - Agent 2] and [Agent 3 - Agent 4]:

$$x^0 = \begin{bmatrix} 30 & 20 \\ 40 & 25 \\ 130 & 80 \\ 140 & 85 \end{bmatrix}$$

In addition, assume $\lambda_i^{t-1} = \lambda_i^0 \left(\frac{1}{1.3}\right)^{t-1}$ for every $t > 1$; that is, the intensity parameter λ_i^{t-1} of the compound Poisson processes decreases, at every communication round $t > 1$, by a factor of 1.3 for every agent i. Then, agents become increasingly self-confident over time and the listening matrix eventually converges to the identity matrix. This might seem reasonable because the abilities to be evaluated at every communication round are increasingly difficult to estimate and it might seem less likely that future information will be received about them. Any other parameter is the same as in the first example; the first listening matrix T^0, the updated beliefs x^1, and the second listening matrix T^1 already show that the initial clusters persist over time:

$$T^0 = \begin{bmatrix} 0.63 & 0.37 & 0 & 0 \\ 0.48 & 0.52 & 0 & 0 \\ 0 & 0 & 0.45 & 0.55 \\ 0 & 0 & 0.52 & 0.48 \end{bmatrix} \quad x^1 = \begin{bmatrix} 33.71 & 21.85 \\ 35.18 & 22.59 \\ 135.52 & 82.76 \\ 134.82 & 82.41 \end{bmatrix}$$

$$T^1 = \begin{bmatrix} 0.47 & 0.52 & 0 & 0 \\ 0.50 & 0.50 & 0 & 0 \\ 0 & 0 & 0.46 & 0.54 \\ 0 & 0 & 0.55 & 0.45 \end{bmatrix}$$

After 20 communication rounds, the two groups are still polarized:

$$x^{20} = \begin{bmatrix} 40.57 & 25.89 \\ 40.57 & 25.89 \\ 135.17 & 82.58 \\ 135.17 & 82.58 \end{bmatrix}$$

Eventually, opinions converge within each cluster (see Fig. 8) while agents become self-confident too rapidly to reach a consensus within the network.

5.3. Third example: two clusters and one not self-confident agent

We now consider the same setting proposed in the previous example except that we set different intensity parameters for the agents with $\lambda_1^{t-1} = 0.8$ and $\lambda_i^{t-1} = 0.1$ for every $i \neq 1$ and for every communication round $t > 1$. Thus, Agent 1 is quite insecure about his ability to evaluate his sources of information, including himself. If any other parameter is unchanged, matrix T^0 will be as in the previous example, while the value of λ_1^1 has a clear effect on the listening matrix T^1 and, therefore, on the updated beliefs x^2:

$$T^1 = \begin{bmatrix} 0.37 & 0.37 & 0.17 & 0.07 \\ 0.50 & 0.49 & 0 & 0 \\ 0 & 0 & 0.46 & 0.54 \\ 0 & 0 & 0.55 & 0.45 \end{bmatrix} \quad x^2 = \begin{bmatrix} 59.27 & 37.09 \\ 34.44 & 22.22 \\ 135.14 & 82.57 \\ 135.20 & 82.60 \end{bmatrix}$$

Agent 1, being quite uncertain about his own estimates of π_i^0, assigns the lowest weight to himself and positive weights to all other agents' opinions. While Agent 3 and Agent 4 maintain their opinions almost unchanged, Agent 1 gets closer to their estimates. As

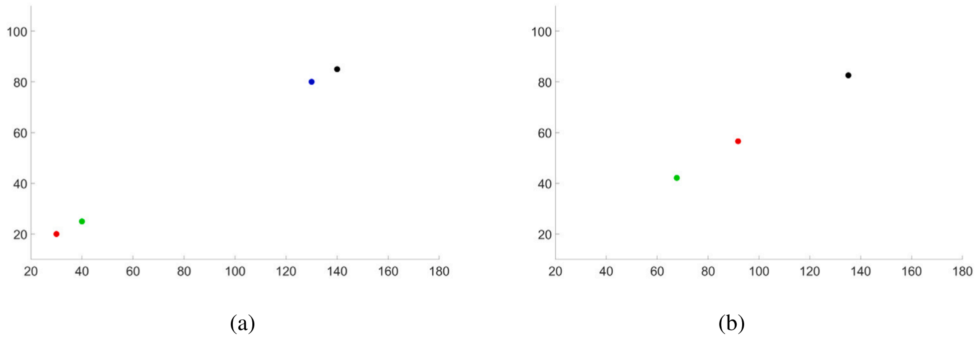

Fig. 9. Third example: (a) agents' initial opinions, (b) agents' opinions after the 10-th iteration.

long as agents reveal their information truthfully, when π_1^1 is shared at the third communication round, Agent 2 believes that Agent 1 is genuinely uncertain about his own reliability in estimating π_1^0. Being close enough to π_2^1, Agent 1's estimate π_1^1 represents a possibly profitable future scenario for Agent 2. Reasonably, Agent 2 still assigns a positive weight to Agent 1 and so, indirectly, to the estimates x_1^2. Therefore, loosely speaking, Agent 1 acts as a bridge between the two groups. Agent 2, although quite self-confident, is willing to assign a positive weight to Agent 1 and thus, albeit indirectly, to Agent 3 and Agent 4.

$$T^2 = \begin{bmatrix} 0.38 & 0.39 & 0.23 & 0 \\ 0.47 & 0.53 & 0 & 0 \\ 0 & 0 & 0.47 & 0.53 \\ 0 & 0 & 0.56 & 0.44 \end{bmatrix}$$

This also happens in the subsequent communication rounds, and the network's opinions eventually converge. Therefore, reaching a consensus between two groups of individuals requires, quite reasonably, at least one agent to be both not self-confident and a trustworthy source of information within his group. A graphical representation of this convergence process is reported in Fig. 9.

5.4. Fourth example: two clusters and one self-confident agent

Let us consider a network with two different clusters [Agent 1- Agent 2 - Agent 3] and [Agent 4] and initial estimates x^0 defined as follows:

$$x^0 = \begin{bmatrix} 30 & 20 \\ 40 & 25 \\ 35 & 22 \\ 140 & 100 \end{bmatrix}$$

In the first communication round, with λ^0 the same as in the first example, the first three agents disregard Agent 4's estimates, while Agent 4 considers the other agents' opinions to be slightly reliable:

$$T^0 = \begin{bmatrix} 0.55 & 0.19 & 0.25 & 0 \\ 0.43 & 0.50 & 0.06 & 0 \\ 0.48 & 0.43 & 0.07 & 0 \\ 0.15 & 0.15 & 0.15 & 0.52 \end{bmatrix} x^1 = \begin{bmatrix} 33.21 & 21.48 \\ 35.36 & 22.65 \\ 34.76 & 22.34 \\ 90.05 & 63.06 \end{bmatrix}$$

For every communication round $t > 1$, set $\lambda_4^{t-1} = 0.05$ and $\lambda_i^{t-1} = 0.7$ for each agent $i \neq 4$. Thus, the agents in the first cluster, unlike Agent 4, are quite unconfident about their own estimates of their abilities to evaluate the unknown parameter. Therefore, in the second communication round, they are willing to assign a significant weight to Agent 4, whose opinion remains unchanged:

$$T^1 = \begin{bmatrix} 0.34 & 0.38 & 0.01 & 0.26 \\ 0.15 & 0.35 & 0.15 & 0.32 \\ 0.18 & 0.37 & 0.15 & 0.28 \\ 0.06 & 0 & 0 & 0.93 \end{bmatrix} x^2 = \begin{bmatrix} 49.15 & 32.98 \\ 52.71 & 35.55 \\ 50.72 & 34.10 \\ 86.49 & 60.45 \end{bmatrix}$$

At the third communication round, the agents in the first cluster, although their opinions are still far apart, assume that Agent 4 is genuinely convinced about her evaluations provided that our setting excludes strategic behavior. This behavior has a clear impact on the other agents, who are, in a way, persuaded by Agent 4 because they are still uncertain about their estimates. The network's opinions converge, in approximately 10 iterations, to Agent 4's opinion; see Fig. 10. This result is perfectly consistent with intuition. To be persuasive, an agent should be (or show herself to be) confident and should face doubtful counterparts.

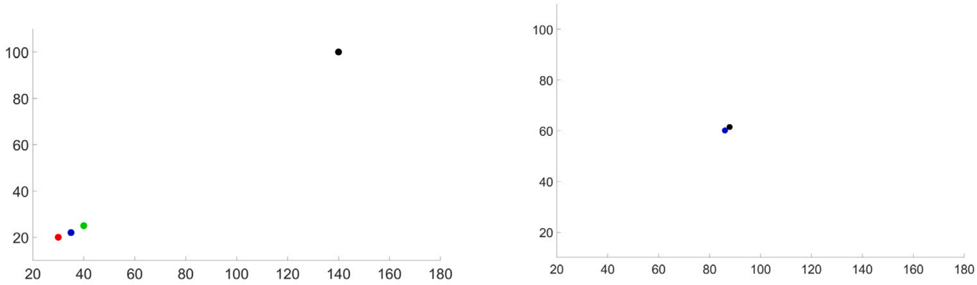

(a) (b)

Fig. 10. Fourth example: (a) agents' initial opinions, (b) agents' opinions after the 10-th iteration.

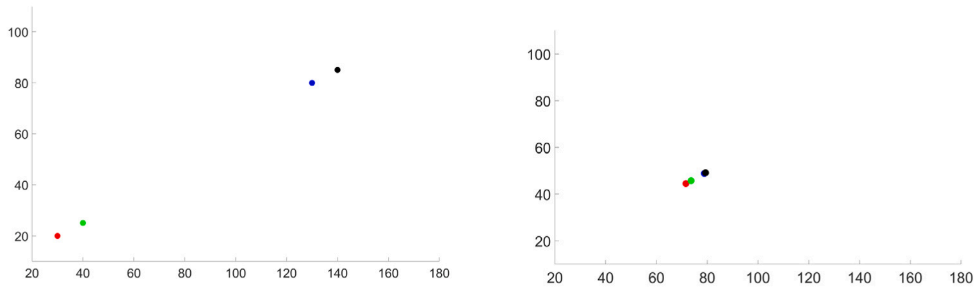

(a) (b)

Fig. 11. Fifth example: (a) agents' initial opinions, (b) agents' opinions after the 3-rd iteration.

5.5. Fifth example: herd behavior

The values of the intensity parameters, appropriately chosen, enable modeling of so-called herd behavior, which is defined as the phenomenon of individuals following others and imitating group behaviors rather than making decisions independently on the basis of their own information. In economic markets, herding behaviors are quite common because new information sources are scarce for most individuals and, thus, they behave similarly to each other in decision-making. Note that having no information is equivalent to having invaluable information. Agents who consider available information to be completely unreliable expect to receive, in the future, new elements or data that will considerably modify their initial estimates. This situation can be mimicked by assuming a high value of λ_i^{t-1} for every agent i at every communication round $t > 1$. Thus, we adopt the setting of the second example with two distinct clusters and we assume $\lambda_i^{t-1} = 0.9$ for $t > 1$. While the first round remains unchanged, now matrix T^1 shows a much faster convergence:

$$T^1 = \begin{bmatrix} 0.36 & 0.36 & 0.26 & 0 \\ 0.44 & 0.30 & 0.11 & 0.13 \\ 0.36 & 0 & 0.36 & 0.27 \\ 0.32 & 0 & 0.35 & 0.32 \end{bmatrix} x^2 = \begin{bmatrix} 61.71 & 38.55 \\ 60.17 & 37.65 \\ 98.55 & 60.66 \\ 101.82 & 62.62 \end{bmatrix}$$

Finally, the network's opinion converges in just three iterations to a consensus. A graphical representation of the convergence is in Fig. 11.

6. Conclusions

Our study proposes a new model of information transmission within a social network that introduces two major innovations. First, multiple communication rounds are justified by the idea that the informative content of the social network has different dimensions, all of which are valuable for the agents. Except for the first interaction, in which the estimates of the unknown parameters are shared, in all the others we consider the exchange of new information about the reliability of the agents. This eliminates the repetition of information, which has always been considered an inevitable feature of an average-based updating process. Second, we suggest a new procedure to determine the value of the weights adopted after each communication round in order to define the updated estimates. We believe that this is also a relevant contribution because it proposes an unexpected bridge between decision theory and portfolio theory exploiting option spreads and portfolio optimization. Interestingly, many practical examples are proposed in which the model predicts reasonable behaviors for the agents. Mainly, the model does not indicate a precise outcome of social interaction. Both consensus and persistent disagreement are permissible outcomes, depending on each agent's self-confidence and sensitivity to estimation errors. The weights assigned to other information sources decrease as agents become more self-confident. Agents who believe their estimates to be correct and reliable rule out any future flow of relevant information that might change their assessments. Therefore, they would give zero weight to any other source of information and would retain their opinion, preventing consensus. Similarly, disagreement will persist if the network is divided into self-confident clusters. In short, achieving consensus relies on the presence of uncertain agents who assign positive weights to alternative sources of information as a hedge against the possibility of unexpected new information.

Declaration of competing interest

None.

Acknowledgments

The research of Sergio Ortobelli Lozza was supported by MIUR-ex60% 2021 sci.resp. Sergio Ortobelli Lozza. The research of Sebastiano Vitali was supported by MIUR-ex60% 2024 sci.resp. Sebastiano Vitali.

Appendix A. Portfolio optimization

We first model the stochastic process that drives the underlying of option strategies. Let $\Pi_i\left(t;\lambda_i^0\right)$ be a Poisson process with intensity parameter λ_i^0 and discrete jump times $t_1, t_2, \ldots$. Construct a new process $\Pi_i^Y\left(t;\lambda_i^0\right)$ by assigning jump Y_1 at time t_1, Y_2 at time t_2, etc., where $Y_1, Y_2, \ldots$ are independent identically distributed (standard) normal random variables.[17] This process can be written as follows:

$$\Pi_i^Y\left(t;\lambda_i^0\right) = \sum_{k=1}^{\Pi_i\left(t;\lambda_i^0\right)} Y_k \tag{A.1}$$

In other words, at time t, Π_i^Y is the sum of $\Pi_i\left(t;\lambda_i^0\right)$ independent identically distributed copies of a random variable Y, where $\Pi_i\left(t;\lambda_i^0\right)$ is a standard Poisson process.[18] Let $\hat{t}$ be the point in time when the decision-maker makes a choice or a consensus is reached. Then agent i's estimated value of the unknown parameter $x_{i,t}^0$ at any time $t > \hat{t}$ is given by the following:

$$x_{i,t}^0 = x_i^0 + \sum_{k=1}^{\Pi_i\left(t;\lambda_i^0\right)} Y_k \tag{A.2}$$

As already mentioned, at the first communication round agent i can invest in a set $\mathcal{B}_i^0$ of N butterflies indexed by i and j, with each butterfly β_{ij}^0 in $\mathcal{B}_i^0$ having a different center strike equal to the estimate x_j^0 of the unknown parameter shared by agent j. Each butterfly β_{ij}^0 is constructed by two long call options with strike $\overline{x}_{ij}^0 = x_j^0 + \delta_i^0$ and $\underline{x}_{ij}^0 = x_j^0 - \delta_i^0$ with $\delta_i^0 \in \mathbb{R}$, respectively, and two short call options with strike x_j^0. Loosely speaking, the value of the parameter δ_i^0 reflects the sensitivity of agent i's utility with respect to a mistake in the estimate of the unknown parameter because it determines both the maximal payoff of the butterfly and the degree of precision required for the payoff to be positive. Let each butterfly consist of call options with maturity τ, with τ being a conceivable point in time long after the decision is made. The random variable $x_{i,\tau}^0$ represents the future value at time τ of the current estimate x_i^0. Thus, we can define the random variable of agent i's discounted payoff $P_{i,\hat{\tau}}\left(\beta_{ij}^0\right)$ of butterfly β_{ij}^0 as follows[19]:

$$P_{i,\hat{\tau}}\left(\beta_{ij}^0\right) = e^{-r_f \cdot \hat{\tau}}\left[\max\left(x_{i,\tau}^0 - \overline{x}_{ij}^0, 0\right) + \max\left(x_{i,\tau}^0 - \underline{x}_{ij}^0, 0\right) - 2\max\left(x_{i,\tau}^0 - x_j^0, 0\right)\right] \tag{A.3}$$

where r_f is the risk-free rate and $\hat{\tau}$ the time to maturity of all options. At the first communication round, each agent i assumes x_i^0 as the initial value of the underlying and constructs a portfolio of butterflies by computing the vector of weights π_i^0 that solve the following optimization problem:

$$\max_{\pi_i^0}\ \gamma_i \mathrm{E}\left[\sum_{j=1}^N \pi_{ij}^0 \cdot P_{i,\hat{\tau}}\left(\beta_{ij}^0\right)\right] - \frac{1}{2}\mathrm{Var}\left[\sum_{j=1}^N \pi_{ij}^0 \cdot P_{i,\hat{\tau}}\left(\beta_{ij}^0\right)\right] \tag{A.4}$$

$$\text{s.t} \sum_{j=1}^N \pi_{ij}^0 = 1$$

[17] We could easily assume a normal distribution with mean zero and time-dependent variance decreasing over time. This would reflect the fact that an agent is unlikely to receive information long after a decision has been made. Alternatively, instead of a standard normal distribution, one could consider a Student's t-distribution allowing for fat tails and thus for a greater chance of extreme values than normal distributions.

[18] A discrete approach is adopted here. This is just a simplifying assumption and the idea can be easily applied in continuous time.

[19] Assume that every option strategy has cost zero. This hypothesis is not only to keep things simple but also because no agent actually buys options.

$$0 \leq \pi^0_{ij} \leq 1 \quad \text{for } \forall i, j$$

where γ_i is the risk tolerance coefficient of agent i, and $\mathrm{E}[\cdot]$ and $\mathrm{Var}[\cdot]$ represent the expected value and the variance. At every time $t > \hat{t}$, the weight $\pi^0_{ij,t}$ assigned by agent i to agent j will be given by the following:

$$\pi^0_{ij,t} = \frac{\max\left(\pi^0_{ij} + \sum_{k=1}^{\Pi_i\left(t;\lambda^1_{ij}\right)} Y_k, \kappa_{ij}\phi\right)}{\sum_{j=1}^{N} \max\left(\pi^0_{ij} + \sum_{k=1}^{\Pi_i\left(t;\lambda^1_{ij}\right)} Y_k, \kappa_{ij}\phi\right)} \tag{A.5}$$

where $\kappa_{ij} \in \{0,1\}$ is an indicator parameter with $\kappa_{ij} = 1$ if and only if $i = j$ and ϕ is a positive scalar arbitrarily close to zero. The ϕ parameter has been introduced to ensure that each agent listens to herself. Thus, the value of the denominator is always positive and $\pi^0_{ij,t}$ is well defined for every pair of agents $\{i, j\}$ at every time t. In addition, the weights sum up to one. We define the random variable of each butterfly's payoff $P_{i,\hat{\tau}}\left(\beta^1_{ij}\right)$ as the sum of the payoffs of N different *sub*-butterflies β^1_{ijh}, each with center π^0_{jh}, with $h = 1, \ldots, N$. In order to compute the payoff of the butterflies, we cannot assume that butterflies are constructed from basket options: given that the weights in every vector π^0_i must sum to one, the equally weighted average of the elements in π^0_i would be identical for every agent *i*. Butterflies' payoffs cannot even be computed by introducing a norm for the distance between the vectors. Because a norm is non-negative by definition, the payoff of every call option would always be non-zero, unless $\pi^0_i = \pi^0_{i,\tau}$.

$$P_{i,\hat{\tau}}\left(\beta^1_{ijh}\right) = e^{-r_f\cdot\hat{\tau}}\left[\max\left(\pi^0_{ih,\tau} - \overline{\pi}^0_{ijh}, 0\right) + \max\left(\pi^0_{ih,\tau} - \underline{\pi}^0_{ijh}, 0\right) - 2\max\left(\pi^0_{ih,\tau} - \pi^0_{jh}, 0\right)\right] \tag{A.6}$$

where $\hat{\tau}$ is the time to maturity of all options. Each *sub*-butterfly β^1_{ijh} is thus constructed by two long call options with strike $\overline{\pi}^0_{ijh} = \pi^0_{jh} + \delta^1_i$ and $\underline{\pi}^0_{ijh} = \pi^0_{jh} - \delta^1_i$, respectively, and two short call options with strike π^0_{jh}. Again, the value of the parameter δ^1_i reflects the sensitivity of agent *i*'s utility with respect to a mistake in the estimate of the unknown parameter.[20] The random variable $\pi^0_{ih,\tau}$ represents the future value at time τ of the current estimate π^0_{ih}. Therefore the random variable of each butterfly's payoff $P_{i,\hat{\tau}}\left(\beta^1_{ij}\right)$, is the sum of the payoffs of N different *sub*-butterflies β^1_{ijh}:

$$P_{i,\hat{\tau}}\left(\beta^1_{ij}\right) = \sum_{h=1}^{N} P_{i,\hat{\tau}}\left(\beta^1_{ijh}\right)$$

Finally, the agent defines the new vector of weights π^1_i solving an optimization problem analog to (A.4).

Appendix B. Convergence

Consider a collection $\mathcal{B}^0_i$ of N butterflies and let Σ be its variance-covariance matrix and μ its N-dimensional vector of expected returns; let e denote the N-vector whose components are all units. The matrix Σ is positive definite. The mean-variance (hereafter MV) problem is as follows:

$$\max_{\pi^0_i}\left\{\gamma_i \mu^\intercal \pi^0_i - \frac{1}{2}\left(\pi^0_i\right)^\intercal \Sigma \pi^0_i | e^\intercal \pi^0_i = 1\right\} \tag{B.1}$$

where γ_i is a scalar parameter, π^0_i is an N-vector of portfolio weights, and $e^\intercal \pi^0_i = 1$ is the budget constraint. The value of the parameter γ_i represents the agent's risk tolerance parameter: the higher the value of γ_i, the more tolerant the investor will be to the risk. The first-order conditions are as follows:

$$\Sigma \pi^0_i + e\psi_i = \gamma_i \mu \tag{B.2}$$

$$e^\intercal \pi^0_i = 1 \tag{B.3}$$

where ψ_i is the Lagrange multiplier for the budget constraint. Then, solving for π^0_i and ψ_i, we have the following:

$$\pi^0_i\left(\gamma_i\right) = \Sigma^{-1}\frac{e}{c} + \gamma_i\left[\Sigma^{-1}\left(\mu - e\frac{a}{c}\right)\right], \tag{B.4}$$

and:

$$\psi_i\left(\gamma_i\right) = -\frac{1}{c} + \gamma_i \frac{a}{c} \tag{B.5}$$

[20] This value could be reasonably assumed to be constant over time.

where efficient set constants a and c are defined as follows:

$$a = e^{\intercal}\Sigma^{-1}\mu \qquad c = e^{\intercal}\Sigma^{-1}e \tag{B.6}$$

Therefore $\pi_i^0(\gamma_i)$ can be expressed as follows:

$$\pi_i^0(\gamma_i) = \alpha_0 + \gamma_i\alpha_1 \tag{B.7}$$

where:

$$\alpha_0 = \Sigma^{-1}\frac{e}{c} \qquad \alpha_1 = \Sigma^{-1}\left(\mu - e\frac{a}{c}\right) \tag{B.8}$$

Equation (B.7) implies that every component $\pi_{ij}^0 = \alpha_{0j} + \gamma_i\alpha_{1j}$ with $j = 1, \dots, N$ of $\pi_i^0(\gamma_i)$ is a linear function of γ_i. If $\alpha_{1j} > 0$ $(\alpha_{1j} < 0)$, then $\pi_{ij}^0(\gamma_i)$ is increasing (decreasing) in γ_i and will be non-negative, provided that $\gamma_i \geq -\frac{\alpha_{0j}}{\alpha_{1j}}$ $(\gamma_i \leq -\frac{\alpha_{0j}}{\alpha_{1j}})$. Let:

$$\gamma_i^l = \max\left\{-\frac{\alpha_{0j}}{\alpha_{1j}} \mid \text{all } j \text{ with } \alpha_{1j} > 0\right\} \tag{B.9}$$

$$\gamma_i^u = \min\left\{-\frac{\alpha_{0j}}{\alpha_{1j}} \mid \text{all } j \text{ with } \alpha_{1j} < 0\right\} \tag{B.10}$$

It follows that:

$$\pi_{ij}^0(\gamma_i) \geq 0 \text{ for all } j \text{ with } \alpha_{1j} > 0 \text{ and for all } \gamma_i \geq \gamma_i^l \tag{B.11}$$

$$\pi_{ij}^0(\gamma_i) \geq 0 \text{ for all } j \text{ with } \alpha_{1j} < 0 \text{ and for all } \gamma_i \leq \gamma_i^u \tag{B.12}$$

If α_{1j} were equal to zero for one or more asset j, in order to have $\pi_{ij}^0(\gamma_i) \geq 0$, the following would be necessary:

$$\alpha_{0j} \geq 0 \text{ for all } j \text{ such that } \alpha_{1j} = 0 \tag{B.13}$$

Then, the condition $\pi_{ij}^0(\gamma_i) \geq 0$ would be verified if and only if $\gamma_i^l \leq \gamma_i \leq \gamma_i^u$ and condition (B.13) were verified. Furthermore, positively weighted minimum variance portfolios lie on a single segment of the minimum variance frontier. Then, consider a compact set Ω of stochastic matrices with strictly positive entries π_{ij}^0. It must be noted that we have to prove the convergence of a left infinite product of matrices with all matrices in Ω, not the convergence of an Ω - Markov chain.[21] Given that Ω is a compact set of Markov matrices,[22] convergence is a natural consequence. This is a classic result present, among others,[23] in Shen (1988).[24]

Appendix C. Time constant weights

Maximum theorem: Let X and Θ be two topological spaces, $f : X \times \Theta \to \mathbb{R}$ be a continuous function on the product $X \times \Theta$, and $C : \Theta \rightrightarrows X$ be a compact-valued correspondence such that $C(\theta) \neq \varnothing$ for all $\theta \in \Theta$.
Define the marginal function (value function) $f^* : \Theta \to \mathbb{R}$ by the following:

$$f^*(\theta) = \sup\{f(x, \theta) : x \in C(\theta)\} \tag{C.1}$$

And define the set of maximizers $C^* : \Theta \rightrightarrows X$ by the following:

$$C^*(\theta) = argsup\{f(x, \theta)\} = \left\{x \in C(\theta) : f(x, \theta) = f^*(\theta)\right\} \tag{C.2}$$

If C is continuous at θ, then f^* is continuous and C^* is upper hemi-continuous with non-empty and compact values. As a consequence, the *sup* may be replaced by *max* and the *arg sup* by *arg max*.
The theorem is typically interpreted as providing conditions for a parametric optimization problem to have continuous solutions with regard to the parameter. In this case, Θ is the parameter space, $f(x, \theta)$ is the function to be maximized, and $C(\theta)$ gives the constraint set over which f is maximized. Then, $f^*(\theta)$ is the maximized value of the function and $C^*(\theta)$ is the set of points that maximize f. The result is that, if the elements of an optimization problem are sufficiently continuous, then some, but not all, of that continuity is preserved in the solutions. In our case, let $f_i : \left(X \subseteq \mathbb{R}^N\right) \times \left(\Theta \subseteq \mathbb{R}^{N\times N}\right) \to \mathbb{R}$ be a continuous risk reward objective function of the optimization problem for agent i. At every communication round $t > 1$, the choice of π_i^{t-1} affects just agent i's payoffs. Because each agent's choice does not affect other agents' payoffs, consider, instead of N distinct optimization problems, a single constraint maximization with the objective function $\sum_{i=1}^N f_i : X^N \times \Theta \to \mathbb{R}$ where $X^N = \Delta^{N\times N}$ is the set of all possible weights to be assigned by each agent i to all the available butterflies being $\Theta = \Delta^{N\times N}$ the Cartesian product of N simplices of dimension N. Recall that each

[21] An Ω - Markov chain is a Markov chain whose transition matrices all belong to Ω.
[22] A Markov matrix is a stochastic matrix with at least one positive column.
[23] The convergence of left products of stochastic matrices has been widely discussed in the literature, such as that by Hajnal (1976) or Seneta (1981).
[24] See proof of theorem 5.1 in Shen (1988).

butterfly has, as underlying, a vector of dimension N (whose entries must sum to one), and we have one butterfly for each of the N agents. Furthermore, each agent's initial estimate coincides with the center strike of one of the available butterflies. The constraint set C over which $\sum_{i=1}^{N} f_i$ is maximized, is a constant function of θ because $C(\theta) = \Delta^{N\times N}$. Then, C is continuous as required by the theorem. Because $\sum_{i=1}^{N} f_i$ is the sum of continuous functions, C^* is an upper hemi-continuous correspondence with non-empty and compact values, and the Kakutani fixed-point theorem applies. If a fixed point were chosen by the agents, both the underlying of butterflies and their weights in the optimal portfolio would be time-independent.

Appendix D. Unidimensionality of opinions

We start with a proof of unidimensional opinions when the transition matrix is constant over time:

Theorem Appendix D.1. *Consider an N-dimensional non-singular, diagonalizable, generic,*[25] *stochastic matrix T. The rank of T^m, when m goes to infinity, is equal to one. Moreover, for m sufficiently large, the rank of T^m approximates to 2, that is, the points represented by the rows of T^m are arranged on a line before converging to a point.*

Proof. Let ξ_k be the k-th eigenvalue of T. Because T is a generic stochastic matrix, we can assume, without loss of generality, that $\xi_1 = 1$ and $\|\xi_k\| > \|\xi_{k+1}\|$ for $k = 1, \dots, (N-1)$ with $0 < \|\xi_k\| < 1$ for $k = 2, \dots, N$. Note that, for every integer m, $\left\|\xi_{k+1}^m\right\|$ is an infinitesimal of higher order than $\left\|\xi_k^m\right\|$ provided that:

$$\lim_{m\to\infty} \frac{\left\|\xi_{k+1}^m\right\|}{\left\|\xi_k^m\right\|} = \lim_{m\to\infty} \frac{\left\|\xi_{k+1}\right\|^m}{\left\|\xi_k\right\|^m} = \lim_{m\to\infty} \left(\frac{\left\|\xi_{k+1}\right\|}{\left\|\xi_k\right\|}\right)^m = 0 \tag{D.1}$$

Then, from a computational point of view, there exists a value M of m sufficiently large such that $\left\|\xi_k^m\right\| \to 0$ for $k = N$ and $\left\|\xi_k^m\right\| > 0$, $\forall k < N$. Recall that, according to the *Rank Nullity Theorem*, for any square matrix A of order N, we have the following:

$$Rank(A) + Dim(Ker(A)) = N \tag{D.2}$$

Because the dimension of the null space corresponds to the geometric multiplicity of the zero eigenvalues, and because a matrix is diagonalizable if and only if the algebraic multiplicity of every eigenvalue equals its geometric multiplicity, we have that $Rank(T^m) \to N-1$ when ξ_N^m goes to zero. Similarly, there exists a sufficiently large value $\tilde{M}$ of m such that $\left\|\xi_k^m\right\| \to 0$ for any $k > 2$ and $\left\|\xi_k^m\right\| > 0$ for $k = 1, 2$. Therefore, for $m = \tilde{M}$, the rank of T^m approximates to 2. Because $\lim_{m\to\infty} \xi_k^m = 0$ for any $k > 1$ we eventually have $lim_{m\to\infty} Rank(T^m) = 1$. □

Now consider a more general case where a specific class of time inhomogeneous Markov chains is considered[26]:

Theorem Appendix D.2. *Let $T_{1:m}$ be any left product of m N-dimensional, non-singular, diagonalizable, generic, stochastic matrices. Assume that for every positive integer m, the matrix $T(m) = T_{1:m}$ is still generic and diagonalizable. Then the rank of $T(m)$, when m goes to infinity, is equal to one. Moreover, for m sufficiently large, the rank of $T(m)$ approximates to 2, that is, the points represented by the rows of $T(m)$ are arranged on a line before converging to a point.*

Proof. Consider the left product $T(m)$ with m any finite positive integer. Because, by assumption, $T(m)$ is diagonalizable for every m, there exists a diagonal matrix $\Gamma(m)$ such that:

$$T(m) = V^c(m)\Gamma(m)V^r(m) \tag{D.3}$$

where $V^c(m)$ is the matrix whose columns are the column eigenvectors of $T(m)$ and $V^r(m)$ is the matrix whose rows are the row eigenvectors of T. Thus $T(m)$ can be expressed as the product of m identical diagonalizable matrices Φ_m:

$$T_m = (\Phi(m))^m \quad \text{with} \quad \Phi(m) = V^c(m)\Gamma(m)^{\frac{1}{m}} V^r(m) \tag{D.4}$$

Matrix $\Phi(m)$ could obviously change over time depending on the matrices that are part of the left product $T_{1:m}$. Because $T(m)$ is not singular, being presumably the product of m non-singular matrices, matrix $\Phi(m)$, while not necessarily stochastic, has one eigenvalue equal to one and all distinct positive eigenvalues with modulus lower than or equal to one for any integer $m \geq 1$. Therefore Theorem Appendix D.1 applies. □

[25] A generic matrix is a matrix whose eigenvalues are all distinct. Non-singularity also ensures that all eigenvalues are different from zero.

[26] This result represents, up to our knowledge, an original contribution to the theory of left convergent matrix products.

Appendix E. Examples

We briefly outline the general setting that characterizes the practical cases outlined in section 5. In every example for any agent $i \in \mathcal{N}$ and for every $l \in \{1,2\}$, we always set $\delta^0_{i,l} = 5$ and $\delta^{t-1}_i = 0.1$ for any $t > 1$.[27] In other terms, the width of any butterfly available after each communication round is identical across both agents and issues so that an estimation error of any of the unknown parameters has an identical effect on each agent's payoff. Conversely, examples differ in both agents' initial estimates x^0_i and intensities λ^{t-1}_i of the compound Poisson processes driving x^0_i and π^{t-1}_i for $i = 1,\dots,N$ and $t \geq 1$.[28] At the first communication round every agent i, given the initial estimate x^0_i of θ and the vector λ^0_i of intensity parameters characterizing the compound Poisson processes driving x^0_i, generates $S = 2000$ equally likely scenarios indexed by $s = 1,\dots,S$. Denote by ${}_s x^0_{i,l,\tau}$ the value of variable $x^0_{i,l}$ at time τ in the s-th scenario. Since θ is a vector, agent i's payoff, at the expiry date $\hat{\tau}$, of each butterfly β^0_{ij} in the s-th scenario is defined as follows:

$$ {}_s P_{i,\hat{\tau}}\left(\beta^0_{ij}\right) = \sum_{l=1,2} e^{-r_f \cdot \hat{\tau}} \Big[\max\left({}_s x^0_{il,\hat{\tau}} - \overline{x}^0_{ijl}, 0\right) + \max\left({}_s x^0_{il,\hat{\tau}} - \underline{x}^0_{ijl}, 0\right) -2\max\left({}_s x^0_{il,\hat{\tau}} - x^0_{jl}, 0\right)\Big] \tag{E.1}$$

where $r_f = 0.01$ is the risk-free rate, $\overline{x}^0_{ijl} = x^0_{jl} + \delta^0_{il}$ and $\underline{x}^0_{ijl} = x^0_{jl} - \delta^0_{il}$ are the strikes of the two long call options, and the time to maturity of all options is set equal to $\hat{\tau} = 20$. Thus, butterfly's payoff is defined as the sum of the payoffs of two distinct butterflies each written on a different underlying. Define agent i's risk-reward function f_i as a combination of the expected value and the variance of the payoffs generated by the portfolio strategy:

$$ \max_{\pi^0_i} \; \gamma_i \, \mathbb{E}\left(\sum_{j=1}^{N} \pi^0_{ij} \cdot {}_s P_{i,\hat{\tau}}\left(\beta^0_{ij}\right)\right) - \frac{1}{2}\mathbb{V}\mathrm{ar}\left(\sum_{j=1}^{N} \pi^0_{ij} \cdot {}_s P_{i,\hat{\tau}}\left(\beta^0_{ij}\right)\right) \tag{E.2}$$

where $\gamma_i = 0,1$ for every i. At each following communication round, the same procedure applies, changing the values of the parameters as required by the specific example.

Data availability

No data was used for the research described in the article.

References

Barone, G., D'Acunto, F., Narciso, G., 2015. Telecracy: testing for channels of persuasion. Am. Econ. J.: Econ. Policy 7 (2), 30–60.
Best, M., Grauer, R., 1992. Positively weighted minimum-variance portfolios and the structure of asset expected returns. J. Financ. Quant. Anal. 27 (4), 513–537.
Blyth, S., 2014. An Introduction to Quantitative Finance. Oxford University Press.
Brandts, J., Giritligil, A.E., Weber, R.A., 2015. An experimental study of persuasion bias and social influence in networks. Eur. Econ. Rev. 80, 214–229.
Cao, M., Morse, A.S., Anderson, B.D., 2008. Reaching a consensus in a dynamically changing environment: convergence rates, measurement delays, and asynchronous events. SIAM J. Control Optim. 47 (2), 601–623.
Corazzini, L., Pavesi, F., Petrovich, B., Stanca, L., 2012. Influential listeners: an experiment on persuasion bias in social networks. Eur. Econ. Rev. 56 (6), 1276–1288.
DeGroot, M.H., 1974. Reaching a consensus. J. Am. Stat. Assoc. 69 (345), 118–121.
DeMarzo, P., Vayanos, J., Zwiebel, J., 2003. Persuasion bias, social influence and uni-dimensional opinions. Q. J. Econ. 118 (3), 909–968.
Egerstedt, M., Martini, S., Cao, M., Camlibel, K., Bicchi, A., 2012. Interacting with networks: how does structure relate to controllability in single-leader, consensus networks? IEEE Control Syst. Mag. 32 (4), 66–73.
Enikolopov, R., Petrova, M., Zhuravskaya, E., 2011. Media and political persuasion: evidence from Russia. Am. Econ. Rev. 101 (7), 3253–3285.
Golub, B., Jackson, M.O., 2010. Naïve learning in social networks and the wisdom of crowds. Am. Econ. J. Microecon. 2 (1), 112–149.
Green, R., Hollifield, B., 1992. When will mean variance efficient portfolios be well diversified? J. Finance 47, 1785–1809.
Hajnal, J., 1976. On products of non-negative matrices. Math. Proc. Camb. Philos. Soc. 79, 521–530.
Hong, H., Kubik, J.D., Stein, J.C., 2004. Social interaction and stock-market participation. J. Finance 59 (1), 137–163.
Hull, J., 2006. Options, Futures, and Other Derivatives, 6. ed., pearson internat. ed edition. Pearson Prentice Hall, Upper Saddle River, NJ.
Lazarsfeld, P.F., Merton, R.K., 1954. Friendship as a social process: a substantive and methodological analysis. Freedom Control Mod. Soc. 18 (1), 18–66.
Markowitz, H., 1952. Portfolio selection. J. Finance 7 (1), 77–91.
Martin, G.J., Yurukoglu, A., 2017. Bias in cable news: persuasion and polarization. Am. Econ. Rev. 107 (9), 2565–2599.
Melguizo, I., 2019. Homophily and the persistence of disagreement. Econ. J. 129 (619), 1400–1424.
Olfati-Saber, R., Murray, R.M., 2004. Consensus problems in networks of agents with switching topology and time-delays. IEEE Trans. Autom. Control 49 (9), 1520–1533.
Olshevsky, A., Tsitsiklis, J.N., 2009. Convergence speed in distributed consensus and averaging. SIAM J. Control Optim. 48 (1), 33–55.
Polanski, A., Vega-Redondo, F., 2023. Homophily and influence. J. Econ. Theory 207, 105576.
Rainer, H., Krause, U., 2002. Opinion dynamics and bounded confidence: models, analysis and simulation. J. Artif. Soc. Soc. Simul. 5 (3).
Rapanos, T., 2023. What makes an opinion leader: expertise vs popularity. Games Econ. Behav. 138, 355–372.
Ren, W., Beard, R.W., 2005. Consensus seeking in multiagent systems under dynamically changing interaction topologies. IEEE Trans. Autom. Control 50 (5), 655–661.

[27] The value of the parameter δ^{t-1}_i is rescaled from 5 to 0.1 for every agent i when $t > 1$ since $\pi^{t-1}_{ij} \in [0,1]$ for every $t \geq 1$.

[28] Recall that the value of λ^{t-1}_i reflects, loosely speaking, the expected amount of future relevant information about the corresponding unknown parameter and, thus, an agent's degree of confidence in her own abilities at evaluating the available sources of information. The lower the quality of agent i's estimate, the higher the value of λ^{t-1}_i, because she expects to receive, with high probability, a lot of new information in the future.

Sarymsakov, T.A., 1961. Inhomogeneous Markov chains. Theory Probab. Appl. 6 (2), 178–185.
Seneta, E., 1981. Non-negative Matrices and Markov Chains, 2nd ed. edition. Springer Series in Statistics. Springer, New York.
Shen, J., 1988. A Geometric Approach to Ergodic Non Homogeneous Markov Chains. CAN Report 98-46. Dept. of Math, UCLA.
Tsitsiklis, J., Bertsekas, D., Athans, M., 1986. Distributed asynchronous deterministic and stochastic gradient optimization algorithms. IEEE Trans. Autom. Control 31 (9), 803–812.
Wolfowitz, J., 1963. Products of indecomposable, aperiodic, stochastic matrices. Proc. Am. Math. Soc. 14 (5), 733–737.