Journal of Economic Behavior and Organization 238 (2025) 107211

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

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

Research paper

(Mis)information diffusion and the financial market[☆]

Tommaso Di Francesco[a,b,*], Daniel Torren-Peraire[c,b]

[a] Center for Nonlinear Dynamics in Economics and Finance (CeNDEF), Department of Quantitative Economics, University of Amsterdam, Roetersstraat 11, NL-1018 WB Amsterdam, The Netherlands
[b] Department of Economics, Ca'Foscari University, Dorsoduro 3246, 30123 Venice, Italy
[c] Institute of Environmental Science and Technology (ICTA-UAB), Universitat Autònoma de Barcelona, 08193 Bellaterra Barcelona, Spain

ARTICLE INFO

JEL classification:
D53
D82
D85
G12
G41

Keywords:
Asset-pricing
Agent-based-model
Information diffusion
Behavioral economics

ABSTRACT

This paper investigates the interplay between information diffusion in social networks and its impact on financial markets with an Agent-Based Model (ABM). Agents receive and exchange information about an observable stochastic component of the dividend process of a risky asset à la Grossman and Stiglitz (1980). A small proportion of the network has access to a private signal about the component, which can be clean (information) or distorted (misinformation). Other agents are uninformed and can receive information only from their peers. All agents are Bayesian, adjusting their beliefs according to the confidence they have in the source of information. We examine, by means of simulations, how information diffuses in the network and provide a framework to account for delayed absorption of shocks, that are not immediately priced as predicted by classical financial models. We investigate the effect of the network topology on the resulting asset price and evaluate under which condition misinformation diffusion can make the market more inefficient.

1. Introduction

Financial markets display a range of empirical regularities that challenge the canonical assumption of a representative agent endowed with Full Information Rational Expectations (FIRE). Asset returns are known to be skewed and leptokurtic, and contrary to the Efficient Market Hypothesis, which posits that prices fully and immediately reflect all available information, there is substantial empirical evidence of frictions and delays in the incorporation of news into asset prices (Huberman and Regev, 2001; Vozlyublennaia, 2014).

A natural starting point for explaining these anomalies is to relax the representative agent assumption. A large body of literature has done so by incorporating heterogeneous agents into otherwise standard asset pricing models. Two major approaches have emerged: the first relaxes rational expectations, showing how *boundedly rational* agents can survive in markets and how their interaction with rational agents generates rich dynamics (Chiarella, 1992; Brock and Hommes, 1997, 1998; Lux, 1998). The second line of work relaxes the full information assumption, focusing instead on how agents acquire and process information in settings with asymmetric or incomplete information (Kyle, 1985; Grossman and Stiglitz, 1980; Barberis et al., 1998). However, existing research has paid less attention to how information is spread across the market.

☆ This work has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 956107, "Economic Policy in Complex Environments (EPOC)".

* Corresponding author at: Center for Nonlinear Dynamics in Economics and Finance (CeNDEF), Department of Quantitative Economics, University of Amsterdam, Roetersstraat 11, NL-1018 WB Amsterdam, The Netherlands.
E-mail address: t.difrancesco@uva.nl (T. Di Francesco).

https://doi.org/10.1016/j.jebo.2025.107211
Received 17 January 2025; Received in revised form 7 July 2025; Accepted 19 August 2025
Available online 27 August 2025

This paper addresses that gap by examining why, in financial markets where agents are highly motivated to find the best available information, many participants still end up following misinformation, which we define as a stochastic process uncorrelated with an asset fundamental value. We argue that the answer lies in how information spreads through social networks and how agents endogenously choose which sources to trust. To investigate this, we create an Agent-Based Model (ABM) where investors are connected via a social network and must form expectations based on both direct and peer-provided signals. While only a subset of agents receives direct information about fundamentals, the others must infer value from their neighbors. Each agent independently decides which information sources to rely on, weighting them according to their perceived historical accuracy. This leads to a dynamic belief-updating process that merges Bayesian learning with a behavioral element. This setup enables us to capture a core empirical tension: even when agents are forward-looking and act optimally given their information sets, misinformation can spread and persist, especially when misinformed agents occupy influential positions in the network. As a result, news absorption into prices is delayed, consistent with empirical findings. Moreover, our model predicts that misinformation shocks tend to diffuse less than true information, in line with the observation that financial markets are not fully efficient but do exhibit some resilience to noise. To validate the model, we calibrate it to match moments of Tesla stock returns, a security often discussed in social media and subject to waves of investor enthusiasm and misinformation. We then use the calibrated model to explore how different network topologies and information structures affect the speed and accuracy of price discovery.

This paper contributes a novel rationale for the persistence of misinformation in financial markets: not as a consequence of irrationality, but as an emergent outcome of optimal behavior under imperfect and socially mediated information.

1.1. Related literature and contribution

Our paper merges insights from two main literatures: one focusing on information diffusion in social networks and one on the effect of information frictions on financial markets. Since the seminal work of Degroot (1974) on consensus formation, scholars have explored multiple belief-updating mechanisms, broadly divided into Bayesian and non-Bayesian classes. In the non-Bayesian tradition, Gale and Kariv (2003) was among the first to embed a network component into the social learning literature, raising the question of whether agents should account for the network structure in decision-making. DeMarzo et al. (2003) argued that information obtained from the network may be biased due to reflexive influences as agents' beliefs may indirectly reflect their own prior beliefs. In our model, this concern is circumvented by having agents update their beliefs only once per period and by assuming that they do not think strategically about the network structure. Acemoglu et al. (2010) extend this literature by analyzing (mis)information diffusion with bilateral updates, while Acemoglu and Ozdaglar (2011) and Kanoria and Tamuz (2013) focus on Bayesian learning under uncertainty. Buechel et al. (2015) emphasize the role of conformity and network centrality in shaping collective wisdom. In contrast, Rusinowska and Taalaibekova (2019) model agents who strategically attempt to influence others. Our framework, by contrast, excludes such strategic behavior and assumes truthful communication, thus differentiating our setting from studies on manipulation through private information access (Benabou and Laroque, 1992). Related to our study, Mueller-Frank and Neri (2013) explore non-Bayesian learning with forecast revisions based on others' actions and Board and Meyer-ter Vehn (2021) study Bayesian learning in adoption settings, where agents observe only one other's behavior, rather than integrating multiple signals as in our model. Finally, recent works have evaluated the use of machine learning in belief formation as in Aymanns et al. (2022), where agents use Deep Q-networks to make binary decisions based on signals received from a network. Unlike our framework, their model does not feature feedback from beliefs to outcomes, which allows agents to filter fake news more easily, though at the cost of accuracy in some situations.

Our contribution lies in proposing a belief-updating mechanism that handles multiple, simultaneous information sources. This mechanism generalizes other social learning models, including Degroot (1974), and is grounded in Bayesian updating with a behavioral component. Specifically, agents construct a time-varying measure of precision for each information source, which they use to derive both the posterior mean and variance of the signal. This second moment is essential in our model, as risk-averse agents use it to determine their optimal asset demands. While some recent approaches, such as Dasaratha et al. (2022), also consider weighted belief aggregation with time-varying weights, they focus only on the mean, whereas the variance plays a central role in our setting.

A complementary strand of literature focuses on the role of information frictions and feedback effects in financial markets. Information frictions have been documented empirically by multiple sources. Huberman and Regev (2001) show that the stock prices of a company, CASI Pharmaceuticals, did not incorporate new information for five months. They point out that the news was initially released as a research article in the journal Nature, but investors reacted only when a Wall Street Journal article reposted the findings of the original study. Behavioral factors can also contribute to information delay. Dellavigna and Pollet (2009) provide evidence of limited attention, demonstrating reduced investor responsiveness on Fridays and identifying profitable strategies that exploit such underreactions. This delay has been studied theoretically, with a focus on insider trading (Kyle, 1985; Benabou and Laroque, 1992; Collin-Dufresne and Fos, 2016) or herding behavior in information acquisition (Banerjee, 1992; Orléan, 1995; Cont and Bouchaud, 2000).

A number of previous studies have examined the role of network dynamics in financial markets. Most of this work has used networks to model imitation among agents, such as in Iori (2002). Panchenko et al. (2013) build on the Brock and Hommes (1998) framework, allowing agents to choose among various trading strategies observed within their network. Khashanah and Alsulaiman (2016) develop an ABM in which agents share their optimal holdings with their neighbors. Wu et al. (2018) also explore different network topologies within an ABM of financial markets, focusing on how traders switch between fundamentalist and chartist strategies. In contrast, Biondo (2020) examines various network structures where agents imitate the discrete trading decisions of their peers, while Bertella et al. (2021) considers scenarios in which agents compare their wealth to that of their neighbors, leading to asset reallocation when their performance lags behind. The novelty of our model lies in the fact that agents do not imitate each other. As already remarked, all agents are forward-looking, and the network is used only to model information flow.

An emerging literature has also studied how positive feedback mechanisms can amplify market inefficiencies. Positive feedback may lead to multiplicity of equilibria in macroeconomic models (Massaro, 2013), as confirmed in experiments (Anufriev and Hommes, 2012; Anufriev et al., 2013; Hommes et al., 2019; Assenza et al., 2021), and in settings such as exchange rate dynamics (De Grauwe and Grimaldi, 2006). In macroeconomic models with behavioral expectations, such as New Keynesian frameworks, positive feedback can destabilize equilibria, especially in multi-country setups (Bertasiute et al., 2020). The persistence of misinformation in markets can also be attributed to such feedback loops as Kopányi-Peuker and Weber (2020) show that even with experience, market participants may not eliminate bubbles or deviations from fundamentals. This suggests that experience alone does not guarantee convergence to rational expectations, aligning with our motivation to study how belief formation under multiple, uncertain signals can lead to persistent inefficiencies.

The rest of the paper is structured as follows: Section 2 introduces the model, focusing on the two blocks that constitutes the ABM, the financial market and the information diffusion process. Section 3 describes the process we use to obtain a realistic calibration of the model. Section 4 presents properties of the model by means of numerical simulations. Section 5 concludes.

2. Model

2.1. The financial market

Consider an economy with I consumers, indexed by $i = (1, 2, \dots, \mathrm{I})$. Consumers are infinitely lived, and at the beginning of every period, they receive the same endowment W_0. They have the same utility over the end of period wealth, given by

$$U(W_t) = -e^{-aW_t}, \tag{1}$$

where $a > 0$ is the coefficient of risk aversion. In order to transfer wealth from the beginning of the period to the end, there are two types of security: a risk-free and a risky asset. We define p_t as the price of the risky asset at a generic time t and normalize the price of the risk-free asset to 1. In every period consumers decide how to allocate their initial endowment, choosing between the two possible securities. At the end of the period they receive profits based on their portfolio, and consume their wealth.[1] Given their participation in the financial market, throughout the paper we use interchangeably the terms consumers, investor and agent. Defining $X_{i,t}$ as the consumer's demand for the risky asset and $M_{i,t}$ as the demand for the safe asset, the allocation choice is subject to the budget constraint

$$p_t X_{i,t} + M_{i,t} = W_0. \tag{2}$$

The risk-free rate is $R > 1$ and the risky asset pays a stochastic payoff which is equal to a dividend claim plus the future price of the asset

$$y_{t+1} = p_{t+1} + d_{t+1}. \tag{3}$$

The presence of the future price on the right-hand side of Eq. (3) ensures a positive feedback mechanism of expectations. This is a well-established feature of financial markets, as documented among others by Heemeijer et al. (2009). There are two stochastic components determining the realization of future dividends

$$d_{t+1} = d + \theta_{t+1} + \varepsilon_{t+1}, \tag{4}$$

with $\varepsilon_{t+1} \sim \mathcal{N}(0, \sigma_\varepsilon^2)$ being pure unobservable noise. The stochastic component θ_{t+1} follows a stationary[2] AR(1) process, $\theta_{t+1} = \beta\theta_t + \eta_{t+1}$ with $\eta_{t+1} \sim \mathcal{N}(0, \sigma_\eta^2)$ and $\beta \in (0, 1)$ and is observable by some agents before its actual realization. One can think of it as information, in a similar spirit to Grossman and Stiglitz (1980) and Gerotto et al. (2019). This implies that the underlying fundamental value of the risky asset is stochastic but that some information about it is revealed in advance. However, this information is not immediately incorporated in the asset price. The end of period t wealth for the ith consumer is given by

$$W_{i,t} = RM_{i,t} + y_{t+1}X_{i,t} = R(W_0 - p_t X_{i,t}) + y_{t+1}X_{i,t}. \tag{5}$$

Agents optimize their end of period wealth, which given the normality of y_{t+1} is also normal. The optimization problem is therefore given by

$$\max_{\{X_{i,t}\}} \left(-exp\left\{ -a\tilde{\mathbb{E}}_{i,t}(W_{i,t}) + \frac{a^2}{2}\tilde{\mathbb{V}}_{i,t}(W_{i,t}) \right\} \right). \tag{6}$$

[1] This assumption allows tractability, as is common in the literature (see, for example, Brock and Hommes (1998)). Within this model, then, wealth does not enter the optimal individual demand function, and does not have an effect on the resulting price. In other words, there are no liquidity restraints as individuals can go long and short on both the risky and riskless assets.

[2] This is a simplifying assumption since it would imply stationary prices. A more accurate representation could be given by considering a growth model for dividends like in Diks and Dindo (2008) and having agents forecast the growth rate. However, while more realistic, this assumption would not change the main insight of the paper, nor affect the calibration which is based on the model returns.

Using Eq. (5) and solving for the optimal choice of risky asset yields

$$X_{i,t} = \frac{\tilde{\mathbb{E}}_{i,t}(y_{t+1}) - Rp_t}{a\tilde{\mathbb{V}}_{i,t}(y_{t+1})}. \tag{7}$$

The notation $\tilde{\mathbb{E}}_{i,t}$ and $\tilde{\mathbb{V}}_{i,t}$ is used to represent subjective expected value and subjective variance for agent i. It is short notation $\mathbb{E}(\cdot|\mathcal{I}_{i,t})$ and $\mathbb{V}(\cdot|\mathcal{I}_{i,t})$, where $\mathcal{I}_{i,t}$ is the information set of agent i at time t, that is before the realization of θ_{t+1}. We set net supply of outside share of the risky asset equal to 0 and use the market clearing condition, imposing supply equal to aggregate demand, to obtain an implicit pricing equation

$$\sum_{i=1}^{I} X_{i,t} = \sum_{i=1}^{I} \left(\frac{\tilde{\mathbb{E}}_{i,t}(y_{t+1}) - Rp_t}{\tilde{\mathbb{V}}_{i,t}(y_{t+1})} \right) = 0. \tag{8}$$

All agents in the model are assumed to be forward-looking in evaluating the asset price. They expect the asset price to be its fundamental value, which is determined by the present discounted value of the stream of future dividends. This implies that in the model there is only a minimal deviation from rationality, given by the asymmetry in information. This is a crucial aspect and one of the main features that distinguish our paper from other prominent works in the literature. Chiarella (1992), Brock and Hommes (1998) and Lux (1998) among others, focus on the coexistence of (rational) fundamentalists and some type of boundedly rational backward-looking agent. The most common type of bounded rationality for the latter is associated with technical trading rules and the literature has usually labeled them as chartists or trend followers (Day and Huang, 1990; Tramontana et al., 2010, 2013; Anufriev et al., 2020; Gardini et al., 2025). This interplay has been shown to generate complex dynamics, and in some cases, even chaotic behavior. A third type of agents which are sometimes included in this framework are sentiment followers (Gardini et al., 2022; Di Francesco and Hommes, 2025) and their presence can contribute to destabilize the market even further. By contrast, in our model, all agents act optimally given the information available and if frictions were removed from the information diffusion process, we would end back in the rational expectations setting. When solving the forward-looking problem, agents in the model assume that other agents behave identically. This assumption leads them to solve the problem as if they were the representative agent. In other words, they consider the aggregate behavior of all agents to be equivalent to their own individual behavior. With this assumption the pricing equation for each individual is given by

$$p_t = R^{-1}\tilde{\mathbb{E}}_t\left(y_{t+1}\right), \tag{9}$$

in which we omit the subscript i and solving by iterating forward, which is done in Appendix A of the Appendix gives

$$p_t = \frac{d}{r} + \frac{\tilde{\mathbb{E}}_t(\theta_{t+1})}{R - \beta}. \tag{10}$$

The first component of Eq. (10) is the usual discounted value of future expected dividends: without information, the price would be constant for all time periods. The second component is specific of our model and imposed by the presence of the observable component of dividend θ. The next section is devoted to describe the mechanism for which agents receive information about this component.

2.2. Actual and perceived law of motion

We now describe the individual beliefs, by distinguishing between the actual law of motion (ALM) of the observable component of dividends θ_{t+1} and the perceived law of motion (PLM) which agents believe to be true. We note that the actual process θ_{t+1} is not publicly available, and it cannot be directly observed even after its realization. A subset of agents, whom we call informed and misinformed, have access to a private signal which they believe to be the true process. They have full confidence in their source of information and as soon as they receive this private signal, they update their beliefs to it. This behavior can be sustained by the fact that they only observe the realization of payoffs. Errors in their payoff forecasts can be attributed to the idiosyncratic noise ε_{t+1} and not necessarily to their source of information. The remaining agents, which we call uninformed, do not have access to the private signal and update their beliefs in a Bayesian way, which we specify below. The ALM is given by the AR(1) process

$$\theta_{t+1} = \beta\theta_t + \eta_{t+1}. \tag{11}$$

This implies that the conditional distribution of θ_{t+1} given the information set $\mathcal{I}_t$ is

$$\theta_{t+1}|\mathcal{I}_t \sim \mathcal{N}\left(\beta\theta_t, \sigma_\eta^2\right). \tag{12}$$

Since we assume that all agents are aware of the structure of the process, we model all prior beliefs as normal distributions. As anticipated we assign agents to three possible categories.

Informed agents. These agents receive the true information θ_{t+1} and base their forecast on it. One can think that this is due to agents having access to privileged or inside information. We prefer to associate this choice with empirical evidence provided by Huberman and Regev (2001) and Peng and Xiong (2006) supporting the idea of different classes of investors with different access to information. Some agents may possess knowledge to process domain-specific information that other, generalist agents, may lack. Their PLM coincides with the ALM. For them we have

$$\theta_{t+1}|\mathcal{I}_t \sim \mathcal{N}\left(\beta\theta_t, \sigma_\eta^2\right). \tag{13}$$

Misinformed agents. These agents think they have perfect foresight like informed agents, but are actually basing their forecast on misinformation. Their PLM is given by

$$\theta_{t+1} = \beta\gamma_t + \sigma_v := \gamma_{t+1}, \tag{14}$$

that is they assume the process follows an AR(1), with same persistence parameter β but different noise term $v_{t+1} \sim \mathcal{N}(0, \sigma_v^2)$. This can be thought of as a misinformation shock. Their conditional distribution is then given by

$$\theta_{t+1}|\mathcal{I}_t \sim \mathcal{N}\left(\beta\gamma_t, \sigma_v^2\right). \tag{15}$$

One can think of these traders as noise traders in De Long et al. (1990) style or akin to sentiment followers. Specifically, in our model, we treat misinformation as a stochastic process that is uncorrelated with actual information and thus with the asset fundamental value. This is different than fake news, which is usually fabricated information that goes against the true information process as in for example Aymanns et al. (2022). One could then ask why these individuals are trading on noise. As Black (1986) puts it, *"One reason is that they like to do it. Another is that there is so much noise around that they do not know they are trading on noise. They think they are trading on information".* The latter explanation is what we argue can sustain this behavior in our model. Although these agents are not questioning if the process they follow is relevant or not, answering this question is not immediate, as, unlike fake news, one cannot ex-post verify the truth in isolation. Since, as we discuss below, misinformation can spread in financial markets, and hence following misinformation can be consistent with the profit-seeking nature of these agents.

Uninformed agents. These agents do not have any private signal, but their PLM coincides with the ALM. The only difference is that they condition their beliefs on their previous period expectation,

$$\theta_{t+1}|\mathcal{I}_t \sim \mathcal{N}\left(\beta\tilde{\mathbb{E}}_{t-1}(\theta_t), \sigma_\eta^2\right). \tag{16}$$

The literature provides ample evidence that not all information is immediately processed by investors upon release. The simplest explanation is that of limited attention. Hirshleifer et al. (2009) finds that investor reactions to earnings announcement are weaker on days when there are multiple simultaneous news releases. On a similar note (Dellavigna and Pollet, 2009) show that investors take more time to process news on Friday. Tetlock (2011) shows that investors overreact to stale information and Gilbert et al. (2012) demonstrate that this causes mispricing in the market by constructing a profitable trading strategy exploiting this finding. More recently Blankespoor et al. (2020) show that there seems to exist some "disclosure processing costs" that would make disclosures not public information, as usually defined, but a form of private information.

2.3. Information diffusion

Agents are socially connected and are organized in a network. Each agent represents a node, and nodes are entirely characterized by beliefs regarding the observable component of dividends θ_{t+1}. The network is static. All edges are exogenously determined and time-invariant. The edges represent the flow of information between nodes. At the beginning of time t agents have normal prior distribution according to their category. They then receive new data in two ways. Informed and misinformed agents observe θ_{t+1} and γ_{t+1} respectively. Moreover they observe them without noise, so that their posterior beliefs are immediately updated to the realizations of the variables[3]

$$\theta_{t+1}|\theta_{t+1} \sim \mathcal{N}\left(\theta_{t+1}, 0\right) \quad \gamma_{t+1}|\gamma_{t+1} \sim \mathcal{N}\left(\gamma_{t+1}, 0\right). \tag{17}$$

Uninformed agents do not receive any private information, but they can receive information from their peers. Formally, for agent i, we assume that they can observe node j prior mean if an edge exists between node i and node j. Based on this information they update their beliefs in a Bayesian way but incorporating a behavioral component. When an agent is connected to another node in the network, we assume that they construct an implicit variance evaluating the forecasting error of the node over time. To do so we compute an exponential moving average of the forecasting error, given as the squared difference between the last observable payoff and the payoff prediction ($\mu_{j,t-1}$) implied by source j, that is

$$EMA_{j,t} = w\left(y_{t-1} - \left(\frac{dR}{r} + \mu_{j,t-1}\right)\right)^2 + (1-w)EMA_{j,t-1}. \tag{18}$$

Then, to map this to a comparable variance of the given source, we multiply the ratio between source j exponential moving average and their own, with the prior variance

$$\sigma_{j,t}^2 = \sigma_\eta^2 \frac{EMA_{j,t}}{EMA_{i,t}}. \tag{19}$$

In simple terms, if they observe that node j has been more accurate then themselves they attach higher confidence to the information received by that node. They will then update their beliefs taking into consideration the information received from their neighbors

[3] This is in a sense an abuse of notation as a Normal distribution with 0 variance is a degenerate distribution with support at the single point θ_{t+1}, known as the Dirac's delta function. However this notation is convenient since it allow us to model the evolution of the beliefs of this category of agents in the general framework.

and the confidence they have in them. An individual will pay more attention to those information sources that they believe are more trustworthy.

Specifically the updating mechanism is behaviorally motivated as follows. At the beginning of the period an agent has their own beliefs about the distribution of θ. They then receive some new information from their neighbors and we argue that a natural way of incorporating this new information is to update their beliefs proportionally to the confidence they have in the source of information. In particular, they update the mean of the distribution by taking a weighted average of their own and their neighbors belief. On the other hand after being provided with more information, they should now be more certain about their belief, hence the variance, measuring the spread of the distribution around their average belief, should shrink. In order to quantify this intuition we rely on Bayesian updating, which offers a method of determining the optimal way in which posterior beliefs should be formed. It is convenient to reparametrize the distribution in terms of precision, rather than variance, so we define $\tau \equiv \frac{1}{\sigma^2}$ and will express normal distributions in the following discussion as $\mathcal{N}(\mu, \tau)$. We then use the following proposition to introduce the updating mechanism.

Proposition 1 (*Bayesian Updating of Beliefs*)**.** *Assume agents have a normal prior distribution with parameters* (μ_0, τ_0) *and receive* K *new information, each with a Normal likelihood* (μ_k, τ_k), $k = 1, 2 \dots K$. *Then agents posterior distribution is Normal, with posterior mean given by:*

$$\mu_P = \frac{\sum_{k=0}^{K} \left(\mu_k \tau_k\right)}{\tau_P} \tag{20}$$

and posterior precision:

$$\tau_P = \sum_{k=0}^{K} \tau_k. \tag{21}$$

Proof. See Appendix D.1.

Eq. (20) shows that the posterior mean is averaging the individual's own prior with the other individual's mean, and the weight is given by the precision of each source of information. Eq. (21) clarifies that posterior precision is the sum of precisions of all sources on information. Intuitively, the more information an agent receives, the more confident they will be about their forecast.

Notice that if the agent has only one connection, this updating procedure is equivalent to using the Kalman filter to filter out the noise in the signal, as we show in Appendix B. In every period t agent i uses this mechanism to update their beliefs and obtain their posterior distribution $\theta_{t+1} \sim \mathcal{N}(\mu_{P,t}, \sigma^2_{P,t})$.

Our mechanism can be seen as a particular case of the naive updating proposed in Golub and Jackson (2010), and is similar to the weighted average mechanism in Dasaratha et al. (2022) but the novelty of our approach is that we simultaneously derive the posterior variance. This object while of no relevance in most information diffusion works, has a crucial role in our work, given the risk averse behavior of our agents. Also different in our case is that the information exchange happens only one time per time step, therefore avoiding any potential bias given by repeated information as is the case in DeMarzo et al. (2003). Moreover the only source of heterogeneity is given by the different categories of agents, but within the same class, beliefs are ex-ante homogeneous. This implies that there are only two innovations or shocks entering the network at each time step. The combination of this property with the autoregressive structure of the component agents are interested in, make so that aggregating over correlated information is unavoidable but not deleterious.

2.4. Payoffs and prices

With the posterior beliefs regarding the observable component of the dividend θ_{t+1}, we can compute beliefs regarding future payoffs. First, the subjective expectations are $\tilde{\mathbb{E}}_t(\theta_{t+1}) = \theta_{t+1}$ for informed agents, $\tilde{\mathbb{E}}_t(\theta_{t+1}) = \gamma_{t+1}$ for misinformed agents, and $\tilde{\mathbb{E}}_t(\theta_{t+1}) = \mu_{P,t}$ for uninformed agents. Using the expression for p_t allows us to compute, (see Appendix C), the conditional expectation:

$$\tilde{\mathbb{E}}_t(y_{t+1}) = \frac{dR}{r} + \frac{R\tilde{\mathbb{E}}_t(\theta_{t+1})}{R-\beta}, \tag{22}$$

and the conditional variance:

$$\tilde{\mathbb{V}}_t(y_{t+1}) = \sigma_\varepsilon^2 + \tilde{\mathbb{V}}_t(\theta_{t+1}) + \frac{\tilde{\mathbb{V}}_t(\tilde{\mathbb{E}}_{t+1}(\theta_{t+2}))}{(R-\beta)^2} + 2\tilde{\mathbb{COV}}_t\left(\theta_{t+1}, \frac{\tilde{\mathbb{E}}_{t+1}(\theta_{t+2})}{R-\beta}\right), \tag{23}$$

which clarifies that heterogeneity in beliefs is completely described by θ_{t+1}. While the expected value of the stochastic payoff is given for every category by Eq. (22), the conditional variance is specific for each category. In particular, starting from Eq. (23), we have:

Informed agents

$$\tilde{\mathbb{V}}_t(y_{t+1}) = \sigma_\varepsilon^2 + \frac{\sigma_\eta^2}{(R-\beta)^2}, \tag{24}$$

since for these agents, θ_{t+1} is not a random variable at time t, and

$$\tilde{\mathbb{V}}_t(\tilde{\mathbb{E}}_{t+1}(\theta_{t+2})) = \tilde{\mathbb{V}}_t(\theta_{t+2}) = \sigma_\eta^2. \tag{25}$$

Misinformed agents

Similarly, for misinformed agents, we have:

$$\tilde{\mathbb{V}}_t(y_{t+1}) = \sigma_\epsilon^2 + \frac{\sigma_v^2}{(R-\beta)^2}. \tag{26}$$

Uninformed agents

Deriving the conditional variance for uninformed agents is more challenging. In particular, the term $\tilde{\mathbb{E}}_{t+1}(\theta_{t+2})$ depends on the network structure and the updating mechanism we have described. This makes the future forecast a random variable, whose distribution depends on the entire network topology, making the derivation analytically intractable. However, since we are deriving the conditional variance of payoffs under the assumption that agents operate as representative agents, we have:

$$\tilde{\mathbb{E}}_{t+1}(\theta_{t+2}) = \beta\tilde{\mathbb{E}}_t(\theta_{t+1}), \tag{27}$$

which is known at time t and, therefore, not a random variable. Under this assumption, we have:

$$\tilde{\mathbb{V}}_t(y_{t+1}) = \sigma_\epsilon^2 + \tilde{\mathbb{V}}_t(\theta_{t+1}). \tag{28}$$

2.5. The role of social learning

In this section, we highlight the importance of our social learning model by comparing the full model to three simpler cases that are naturally nested within our framework.

The first case is that of full-information, which in our setting obtains if the proportion of informed agents is set to 1. In this case using Eq. (10) the resulting price is

$$p_t = \frac{d}{r} + \frac{\theta_{t+1}}{R-\beta}. \tag{29}$$

The second case features traders' heterogeneity but no social learning. In this framework, agents keep their expectations constant at their prior beliefs. Defining I as the total number of agents, and λ and ξ as the proportions of informed and misinformed agents respectively, and imposing a prior mean of 0 for uninformed agents, the resulting price evolves according to:

$$p_t = \frac{d}{r} + \frac{\lambda\theta_{t+1} + \xi\gamma_{t+1}}{R-\beta}. \tag{30}$$

The third case modifies the previous baseline by adding perhaps the most notable social learning mechanism, that of Degroot (1974), which occurs when uninformed agents assign equal weight to all sources of information available. For this case, we consider a fully connected network, allowing us to derive analytical results. In this framework, every uninformed agent receives λI signals from informed agents, ξI signals from misinformed agents, and $(1-\lambda-\xi)I$ signals from uninformed agents (including themselves).

Moreover, every uninformed agent has identical beliefs. In particular, assuming that the homogeneous variance attached to every piece of information is $\sigma^2 > 0$, we have for uninformed agents:

$$\tilde{\mathbb{V}}_t(\theta_{t+1}) = \frac{\sigma^2}{I}. \tag{31}$$

It is worth noting that the posterior variance approaches 0 as the network size approaches infinity. More precisely, the driving factor here is the number of signals each agent receives, which equals the total number of agents.[4] This result is general, as formalized in the following proposition:

Proposition 2 (*Vanishing of Posterior Variance*). *Assuming that the variance associated with each source of information is positive and bounded, then:*

$$\lim_{K\to\infty} \tilde{\mathbb{V}}_t(\theta_{t+1}) = 0. \tag{32}$$

Proof. See Appendix D.2.

This proposition implies that uninformed agents become increasingly confident in their posterior beliefs as they receive more signals, eventually acting as if their expectations are certain. This feature results in a higher quantity demanded in the market compared to the baseline case, leading to increased volume and more volatile returns. However, the symmetry in the updating process ensures that the return distribution remains symmetric.

The posterior mean for uninformed agents is given by:

$$\tilde{\mathbb{E}}_t(\theta_{t+1}) = \lambda\theta_{t+1} + \xi\gamma_{t+1} + (1-\lambda-\xi)\tilde{\mathbb{E}}_{t-1}(\theta_t). \tag{33}$$

[4] Since their own beliefs are included.

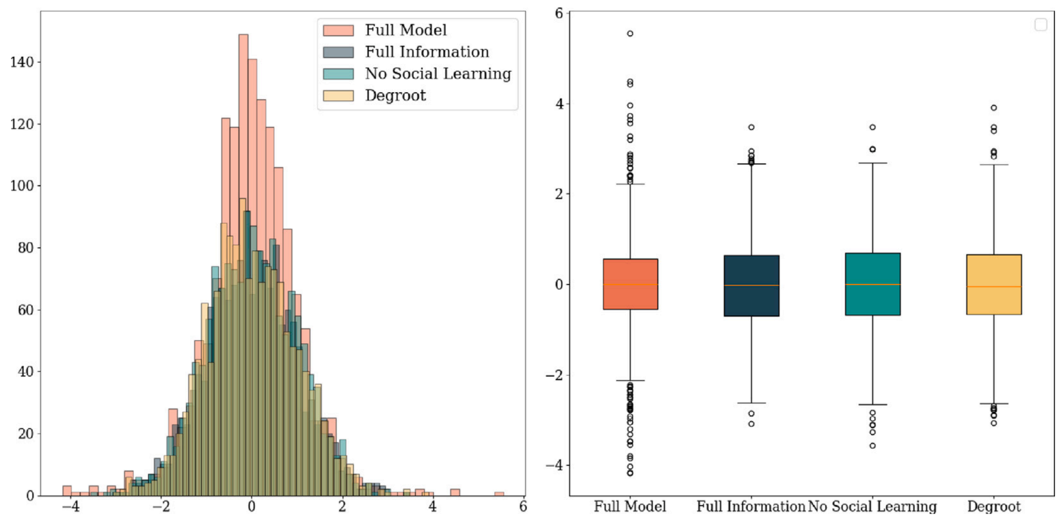

Fig. 1. Return distributions for the three models.

The resulting pricing equation is then:

$$p_t = \frac{d}{r} + \left(\lambda \frac{\frac{\theta_{t+1}}{R-\beta}}{V_I} + \xi \frac{\frac{\gamma_{t+1}}{R-\beta}}{V_M} + (1-\lambda-\xi)\frac{\frac{\tilde{\mathbb{E}}_t(\theta_{t+1})}{R-\beta}}{V_U} \right) \left(\frac{\lambda}{V_I} + \frac{\xi}{V_M} + \frac{(1-\lambda-\xi)}{V_U} \right)^{-1}, \tag{34}$$

with:

$$V_I := \sigma_\varepsilon^2 + \frac{\sigma_\eta^2}{(R-\beta)^2}, \quad V_M := \sigma_\varepsilon^2 + \frac{\sigma_v^2}{(R-\beta)^2}, \quad V_U := \sigma_\varepsilon^2 + \frac{\sigma_\eta^2}{I}.$$

Rearranging, this reads as:

$$p_t = \frac{d}{r} + \frac{1}{R-\beta} \frac{\lambda \left(\theta_{t+1}\right) V_M V_U + \xi(\gamma_{t+1}) V_I V_U + (1-\lambda-\xi)\tilde{\mathbb{E}}_t(\theta_{t+1}) V_I V_M}{\lambda V_M V_U + \xi V_I V_U + (1-\lambda-\xi) V_I V_M}. \tag{35}$$

We now compare the three models to the full model, in which agents are connected in a small-world network. To ensure comparability, we simulate the three models for 30 different realizations of the stochastic processes using different seeds and show the distribution of the normalized model-implied returns in Fig. 1. The parameters used are the same for all models and are fixed to the values in Table 1.

We compare the return distributions for all models and plot a histogram of the distribution and boxplots for each return series. The full-information case, the baseline with no social learning and the DeGroot model all exhibit close to normal return distributions and no excess kurtosis. The introduction of our social learning mechanism is not only more realistic but also necessary to match empirical evidence.

In the next section, we turn to the calibration of the full model and the exploration of its properties.

3. Empirical calibration

First, we summarize the model by offering a more accessible visualization of the sequence of events taking place in each time step t

1. Agents update their prior beliefs regarding the observable component of dividends θ_{t+1}, according to their own category;
2. Agents compute the average forecast error over payoffs of the nodes they are connected to. The last payoff considered by them in the computation is y_{t-1};
3. Agents use the updating mechanism to obtain their posterior beliefs regarding the observable component of dividends θ_{t+1};
4. The implied expected payoff y_{t+1} is computed by each agent;
5. The resulting price p_t and individual demands are computed.

Then we discipline the model and its many parameters by calibrating it with the following strategy. We target the daily returns of the TESLA stock from September 2023 to September 2024. In the literature a common choice to obtain a reliable calibration is to use some index, typically the S&P 500 index (Schmitt and Westerhoff, 2017). In this context we argue that it is more natural to work with a single stock and in particular with one which features a high percentage of retail investors participation. In TESLA case the ownership attributed to the public is around 40%. Moreover individuals and insiders make up 12% of the total ownership, thus making the stock an ideal candidate to study the effects of information frictions. The goal of the empirical calibration is not to exactly match the data-generating process of a specific asset, but rather to discipline key parameters and ensure that the model generates realistic dynamics. Nonetheless we use Appendix E to assess robustness by repeating the calibration exercise for all stocks in the S&P 500, ans confirm that the dynamics we investigate below are not the byproduct of the choice of this particular asset. We then deal with the calibration of the model by following a three-step procedure.

First, some parameters are specific to the network topology. The one we use as our benchmark is the Watts and Strogatz (1998) Small World Network. The network is generated by starting with a regular lattice of a given degree, assigning each node to one of the three categories introduced in Section 2.3 according to specified proportions. Finally we rewire a fraction of edges randomly by a given probability of rewiring, introducing long-distance connections in the network. The main features of such a topology are local clustering, short-average path length and almost homogeneous degree of connection across nodes.

Second we keep some parameters constant at values which are arbitrarily chosen after experimentation with the model, as they do not affect the model dynamics. Total time steps are sufficient to ensure that behaviors driven by initial conditions are absorbed in the long run. The number of agents in the model is of relative importance only when combined with network specific parameters. The gross risk free rate is the daily equivalent corresponding to the average risk free rate in the period September 2023-September 2024 as proxied by the Market Yield on U.S. Treasury Securities at 1-yr Constant Maturity. The constant component of dividends[5] is calibrated to match the daily closing price at the beginning of the period: $d = p_0 \cdot r$. The coefficient of constant risk aversion is at a value common in the literature, see for example Chetty (2006). Third, for some parameters we explore their effect on the model for a wide range of values by Sobol sensitivity analysis (Sobol, 2001). Sobol sensitivity analysis is a global variance-based method used to quantify the contribution of individual input variables, or in our case parameters, to the output variance of a model. Given a model of the form $Y = f(X_1, X_2, \ldots, X_k)$, where Y is a scalar output, the first-order effect of a factor X_i can be expressed as

$$V_{X_i}\left(\mathbb{E}_{X_{\sim i}}(Y|X_i)\right), \tag{36}$$

where X_i represents the ith factor, and $X_{\sim i}$ denotes the matrix of all factors except X_i. The inner expectation operator $\mathbb{E}_{X_{\sim i}}(Y|X_i)$ represents the mean of Y over all possible values of $X_{\sim i}$ while keeping X_i fixed. The outer variance operator is taken over all possible values of X_i. The associated sensitivity measure, known as the first-order sensitivity index, is then given by:

$$S_i = \frac{V_{X_i}\left(\mathbb{E}_{X_{\sim i}}(Y|X_i)\right)}{V(Y)}. \tag{37}$$

S_i is normalized, as $V_{X_i}\left(\mathbb{E}_{X_{\sim i}}(Y|X_i)\right)$ ranges between zero and $V(Y)$. $V_{X_i}\left(\mathbb{E}_{X_{\sim i}}(Y|X_i)\right)$ measures the first-order (additive) effect of X_i on the model output, while $\mathbb{E}_{X_i}\left(V_{X_{\sim i}}(Y|X_i)\right)$ is customarily referred to as the residual.

Another commonly used variance-based sensitivity measure is the total effect index, defined as:

$$S_{T_i} = \frac{\mathbb{E}_{X_{\sim i}}\left(V_{X_i}(Y|X_{\sim i})\right)}{V(Y)} = 1 - \frac{V_{X_{\sim i}}\left(\mathbb{E}_{X_i}(Y|X_{\sim i})\right)}{V(Y)}. \tag{38}$$

The total effect index S_{T_i} measures the total contribution of X_i, which includes both the first-order and higher-order effects (interactions) of the factor. One way to interpret this is by considering that $V_{X_{\sim i}}\left(\mathbb{E}_{X_i}(Y|X_{\sim i})\right)$ represents the first-order effect of $X_{\sim i}$, so that $V(Y)$ minus $V_{X_{\sim i}}\left(\mathbb{E}_{X_i}(Y|X_{\sim i})\right)$ must give the contribution of all terms in the variance decomposition that include X_i.

In Fig. 2 we report the total order Sobol index for the seven parameters in the y-axis, computed on three outputs of the model, namely: price variance in excess of a model with a representative fully informed investor which we use as benchmark, skewness and kurtosis of the model returns. All results are obtained from using 128 samples for each of the 9 parameters of interest, with 30 different seeds per combination, resulting in a total of 34 560 simulations.

Excess variance seems to be driven by two parameters, β and σ_η, which combine to determine the unconditional variance of the observable component of dividends. The skewness of the model returns is driven by the proportions of informed and misinformed agents, λ and ξ, and by the size of the unobservable noise component of dividends σ_ε. A similar result is obtained for the kurtosis of the model returns, although now the proportion of misinformed agents seems to play a smaller role. The number of informed and misinformed agents in the model however has both a direct and indirect effect. The former is due to the fact that these agents are also actively operating in the market and therefore their beliefs are directly reflected in the price. The latter is due to the effect that their communication has on the beliefs of uninformed agents, which in turn affect the price. Since we are mainly interested in this last effect we decide to keep the proportion of informed and misinformed agents at a constant value. The value we chose is 5% of the population for both classed, in order to have a reasonable number of agents in each category while still being able to attribute the majority of the price dynamics to the uninformed agents. The volatilities of the unobservable and observable component of dividends combine to determine the noise-to-signal ratio of the information. We therefore normalize the value of σ_ε to 1. Finally, as

[5] Although TESLA does not currently pay dividends, one could interpret this as an expected average dividend based on expected earnings and an expected payout ratio. Mathematically this value has the only effect of shifting prices and has no impact on the dynamics.

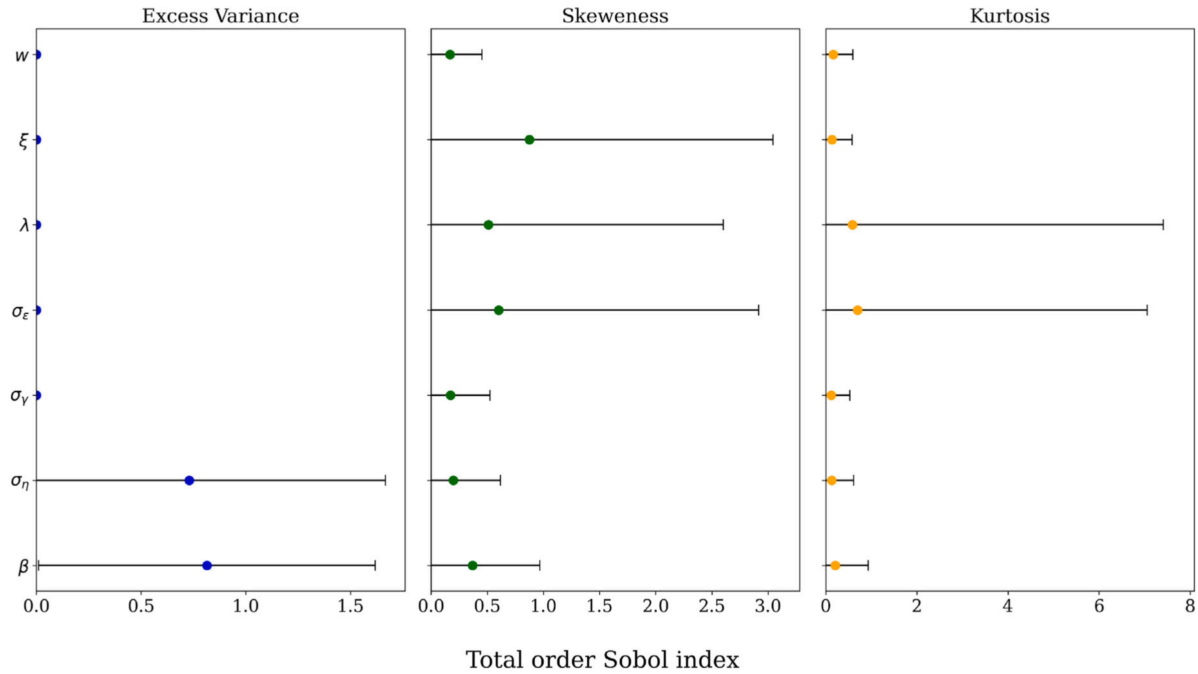

Fig. 2. Sobol sensitivity analysis.

already mentioned, β and σ_η jointly determines the variance of the information. We therefore fix one of them, namely $\beta = 0.5$ and use the other to vary the variance of the information. This leaves us with two parameters to calibrate σ_η and σ_v and we do so by using two moments of the TESLA stock returns: skewness and kurtosis. For this we use the Sequential Neural Posterior Estimation (SNPE) proposed by Papamakarios and Murray (2016) and popularized to economics ABMs by the recent work of Dyer et al. (2024). SNPE is a likelihood-free inference method that uses neural networks to directly approximate the posterior distribution of model parameters.[6] The key idea behind this method is to simulate data from the model at different parameter values and use a neural network to learn the conditional density $q_\phi(\theta|x)$, where θ are the parameters and x is the observed data (e.g., TESLA stock skewness and kurtosis).

The process works as follows:

1. **Simulate Data**: For a given set of parameters θ_n, simulate data x_n from the ABM.
2. **Neural Network Training**: Use the simulated parameter-data pairs (θ_n, x_n) to train a Mixture Density Network (MDN). The MDN outputs a conditional probability distribution $q_\phi(\theta|x)$, which serves as an approximation to the true posterior $p(\theta|x)$.
3. **Posterior Approximation**: After training, the neural network provides an approximation of the posterior distribution for any observed data x_o, by maximizing the likelihood

$$\max_{\phi} \frac{1}{N} \sum_{n=1}^{N} \log q_\phi(\theta_n | x_n). \tag{39}$$

4. **Sequential Refinement**: The method iteratively refines the prior distribution $p(\theta)$ based on the learned posterior, improving the efficiency of the simulations by focusing on plausible regions of the parameter space.

We refer the reader to Dyer et al. (2024) for an in depth evaluation of the method. Other methods such the Simulated Method of Moments (SMM) as in Franke and Westerhoff (2012, 2016) could be used, but SNPE offers two main advantages. While SMM estimates parameters by minimizing the distance between empirical and simulated moments, SNPE leverages neural networks to approximate the full posterior distribution of the parameters. This enables Bayesian estimation and allows the researcher to incorporate prior information. Most importantly for our purpose, SMM must be rerun for each new dataset, while SNPE can be amortized, that is, it allows for fast re-estimation across different observations after a single training phase. While in our main analysis we only use data for TESLA, in our robustness analysis we recalibrate the model to all assets in the S&P500 and SNPE allowed us to obtain the posterior parameters efficiently. Moreover we assess the ability of the method to recover the true posterior

6 This is in contrast with classical Bayesian methods in which the posterior is computed as the product of the likelihood and the prior.

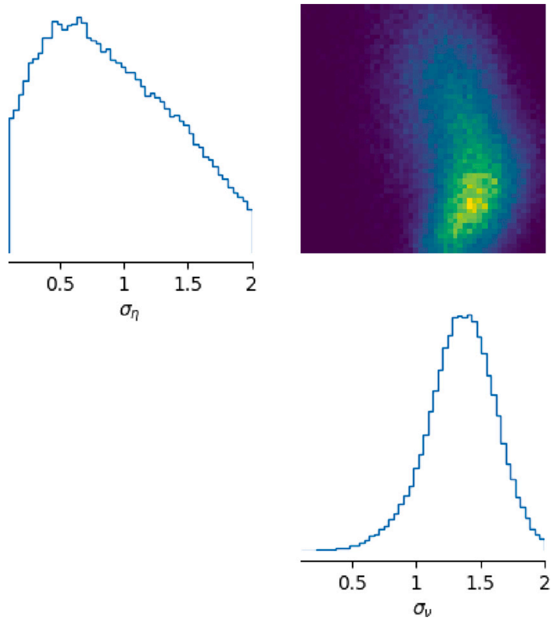

Fig. 3. Posterior distribution.

Table 1
Parameter calibration.

Calibrated parameters				
Parameter	Range explored	Value	Source	Description
T	–	1500	Own calibration	Total time steps
I	–	200	Own calibration	Number of agents
R	–	1,0001	Average daily risk free rate 09:2023–09:2024	Gross risk free rate
d	–	0,021	Directly calibrated	Constant component of dividends
a	–	1	Chetty (2006)	Coefficient of constant risk aversion
σ_ε	$[0.0001, 2]$	1	Normalized	Std of the unobservable component of dividends
w	$[0.1, 0.9]$	0.9	Own calibration	Memory parameter of the exponential moving average
λ	$[0.01, 0.3]$	0.05	Own calibration	Proportion of informed agents
ξ	$[0.01, 0.3]$	0.05	Own calibration	Proportion of misinformed agents
Estimated parameters				
Parameter	Prior	Posterior		Description
		Mean	Median	Std	
σ_η	Uniform (0,2)	0.89	0.84	0.47	Std of information shocks
σ_ν	Uniform (0,2)	1.34	1.36	0.26	Std of misinformation shocks

by using a synthetic dataset generated by the model in Appendix F, which serves to show that in our setting the neural network approximation is accurate.

For the real data estimation we take an agnostic position on the prior. Given that we are estimating variances we use a uniform prior distribution, with support $[0.1, 2]$ and report the results as well as the other parameter values in Table 1. The posterior distribution is extremely well behaved and uni-modal as we can report in Fig. 3, with median values $\sigma_\eta = 0.84$ and $\sigma_\nu = 1.36$. It is interesting to notice that shocks to misinformation are slightly larger than shocks to information. This is consistent with misinformation being more sensational than actual information, in order to generate more attention.

4. Numerical simulations

With our calibration we can then show key features of the model. We do so in Fig. 4 by using the average over 30 Montecarlo simulations with different stochastic seeds. In panel (a) we plot the network structure used in the simulations. Agents position is randomly determined in the beginning on the experiment and then kept fixed, in order to allow us to track each agent evolution in the different simulations. Panel (b) shows a scatter plot of cumulative profits and average forecast error for each agent. The accuracy of uninformed agents lies between that of the other two categories. In this configuration, misinformation does not appear

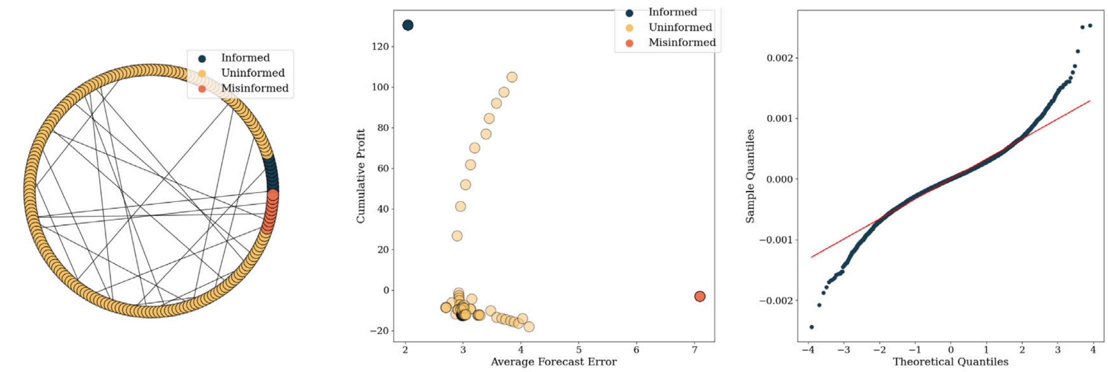

(a) Network Structure (b) Cumulative Profits (c) QQ Plot

Fig. 4. Small world Network.
Panel (a): Small World Network structure. Each dot represents an agent, edges represents connections. Panel (b): average forecast error and profit for each agent. Panel (c): model based quantiles (blue) vs. theoretical (red) quantiles based on a Normal distribution with same mean and variance. Note: Data are obtained from 30 simulations with different stochastic seeds. **Network parameters** are: Network density = 0.1, Probability of rewiring = 0.01. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

to spread significantly into the network. Although there is a generally linear positive relationship between accuracy and profits, there are some notable exceptions. While most uninformed investors incur small gains or even losses, a small subset earns large profits, sometimes almost as high as those of informed agents. The reason for this pattern is that these agents are directly linked to a source of information. They receive information relatively early and act on before others, allowing them to profit the most from the delayed price adjustment caused by the gradual spread of information in the network. Moreover as already pointed out, uninformed agents with many connections might have a lower posterior variance than informed agents, therefore taking larger positions and profiting more from the price movement in the direction of their beliefs. Lastly we look at the impact on returns. In panel (c) we use a QQ plot to identify the presence of fat tails. We compare the simulated quantiles with the theoretical quantiles from a Normal distribution with the same mean and standard deviation of model's returns. In Table 4 we report the summary statistics of the model returns. Skewness and kurtosis are the target moments we have used for the calibration and are 0.04 and 5.50 in the simulated data and 0.25[7] and 5.54 in the empirical counterpart. We also report the profit (or loss) incurred by the misinformed agents in the model. In this case it is negative, showcasing that while misinformation is spreading in the network it does not diffuse outside of the local cluster of connections to misinformed agents.

We can then analyze the implications of the model for the absorption of information in the market using an impulse response function (IRF) analysis. In our framework, shocks propagate through two distinct channels. First, there is a *direct channel* as the shock directly influences expectations, affecting investors' demands and ultimately feeding into market prices. Second, there is an *indirect channel* as the resulting prices and communication between agents feed back into how uninformed traders update their weighting matrices, which determine the relative importance they assign to different sources of information.

If we were to introduce a shock to the θ process at time τ, we could not attribute the resulting effects only to this specific shock because earlier shocks would still influence both the network's structure and the ongoing diffusion of information. To isolate the causal impact of a single shock, we proceed as follows. We begin by simulating the model for a sufficiently long period to ensure that the weighting matrix converges to a stable representative state. We then record this matrix and use it to initialize a new market scenario in which no further updates to the weighting matrix occur. By suppressing all noise until time $\tau = 10$, agents have no reason to alter their weightings, which therefore remain fixed at the initial distribution. Under these controlled conditions, we introduce a one-standard-deviation shock to either θ (an information shock) or γ (a misinformation shock). This setup allows us to compute the resulting impulse response and entirely capture the direct effect of the shock. Fig. 5 shows the impulse response of the price to a one-standard-deviation shock.

The results align with patterns observed in empirical studies. While misinformation shocks do propagate into market prices, their overall impact is weaker than that of information shocks. In line with the findings of Clarke et al. (2021), the market appears able to filter out some of the noise generated by misinformation. By contrast, information shocks not only exert a stronger influence on prices but also continue to shape price dynamics over time. In both cases, the delayed peak of the IRF is consistent with the assumption that information diffuses gradually through the network and only reaches its full impact on prices once uninformed agents have processed it (Huberman and Regev, 2001). This finding also offers a rationale for the evidence provided by the literature about the existence of a momentum factor in stock markets (Jegadeesh and Titman, 1993), since for as long as information is not fully diffused in the network, past returns may predict future returns.

[7] While normally one would expect negative skewness in financial markets, the period we considered was associated with a series of favorable earnings reports and positive growth expectations.

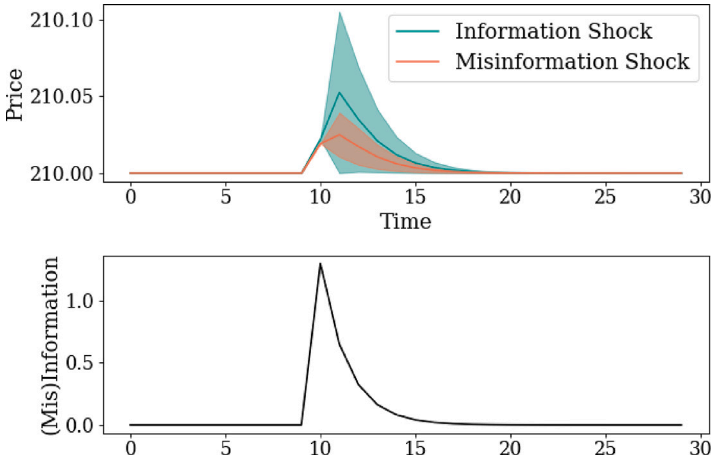

Fig. 5. Impulse response to a (mis)information shock at time $\tau = 10$.

Table 2
Parameters of the Stochastic Block Network.

Parameter	Value
Number of partitions	2
Density of intra-groups edges	0.1
Density of inter-groups edges	0.001

4.1. The effect of different network topologies

The use of social networks has drastically reshaped communication in recent years. The network topology plays a crucial role in enabling rapid information sharing but also raises concerns about how beliefs are formed. In recent years, social media platforms have shifted from a friend-based structure, where connections are reciprocal, to a follower-based structure, where a few influencer nodes can have a disproportionately large number of connections. Moreover, the way many algorithms are designed to maximize user engagement can lead to the creation of echo chambers, where individuals are exposed only to information that confirms their existing beliefs. While the small-world network structure we have used so far is a good approximation of the friend-based structure, we now explore the impact of polarization and follower-based structures on the model dynamics by adopting different network topologies. Crucially, to isolate the effect of the network structure, we repeat the analysis while keeping the variable parameters fixed.

Stochastic Block Network

The first scenario we analyze is a completely polarized society. We create it by using a Stochastic Block Model (Holland et al., 1983), in which we partition the nodes in order to create two clusters. Informed and misinformed agents are separated and assigned to either the information block or the misinformation one. We denote by density of intra-groups edges, the likelihood of having an edge between agents belonging to the same cluster. This is higher than the density of inter-groups edges, regulating connections between agents belonging to different blocks. Thus not only the network is partitioned, but communication between the two groups is scarce. We report the network specific parameters in Table 2.

As before we simulate the model 30 times with different stochastic seeds. Results are displayed in Fig. 6.

Panel (a) illustrates the network structure, which now consists of two clusters each exhibiting small-world architectures. Under these conditions, uninformed agents are less likely of receiving reliable information and misinformation simultaneously. The consequences of this segregation are clearly visible in Panel (b), where agents in the information block earn higher profits than those in the misinformation block. Panel (c) shows that, although the overall shape of the returns distribution remains qualitatively similar, the leptokurtosis is now more pronounced. Misinformed agents continue to incur losses, though these losses are smaller than in the baseline scenario. While these agents can influence their own cluster, their erroneous beliefs fail to fully penetrate the market's pricing mechanism. As a result, uninformed agents situated in the informed cluster capitalize on their disadvantaged counterparts in the misinformed segment of the network.

Scale Free Network Until now we have simulated societies in which agents have equal opportunities of sharing their beliefs, given that the average degree was rather homogeneous across the network. For our next scenarios we opt to work with societies in which certain individuals can have a disproportionate impact on the network.[8] This concept is similar to the one of "guru", discussed

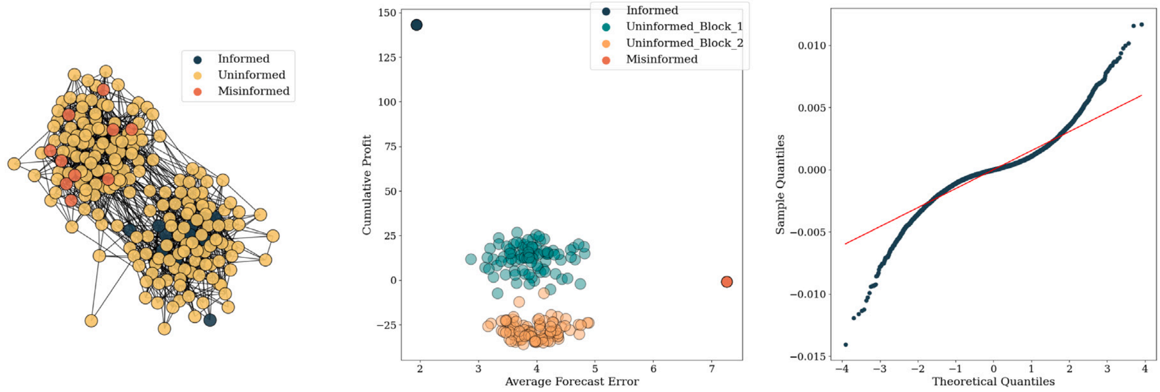

(a) Network Structure (b) Cumulative Profits (c) QQ Plot

Fig. 6. Stochastic block network.
Panel (a): Stochastic Block Network structure. Each dot represents an agent, edges represents connections. Panel (b): average forecast error and profit for each agent. Panel (c): model based quantiles (blue) vs. theoretical (red) quantiles based on a Normal distribution with same mean and variance. Note: Data are obtained from 30 simulations with different stochastic seeds. **Network parameters** are: Number of Partitions = 2, Density of Intra-Groups Edges = 0.1, Density of Inter-Groups Edges = 0.001. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3
Parameters of the Scale Free Network.

Parameter	Range
Alpha	0.41
Beta	0.54
Gamma	0.05

Note: Alpha is the probability for adding a new node connected to an existing node chosen randomly according to the in-degree distribution. Beta is the probability for adding an edge between two existing nodes. Gamma is the probability for adding a new node connected to an existing node chosen randomly according to the out-degree distribution. For a detailed explanation of the parameters role we refer to page 2 of Benabou and Laroque (1992).

in Tedeschi et al. (2012). Gurus are agents that are most imitated by others and can emerge endogenously in the market. The main difference is that in their model, edges are created by a mechanism of preferential attachment based on wealth. Instead we use an exogenous mechanism so by creating a directed[9] Scale-Free Network (Bollobás et al., 2003) and forcefully allocating either informed or misinformed agents in the nodes with most outward connections. The network specific parameters are reported in Table 3

We begin by analyzing the case where informed agents are the most central. The network topology is displayed in panel (a) of Fig. 7. Panel (b) highlights that informed agents benefit significantly from their prominent position in the network, earning much more than in previous scenarios. This is because they can directly reach most of the uninformed traders. Since communication is lagged, informed agents push the price in the direction in which they have already taken a position. This results in the majority of uninformed agents incurring small losses. As everyone receives the information almost simultaneously, there are no additional gains to be made; in fact, uninformed agents lose money on average to informed agents who acted earlier based on their private information. Even so, it is still better for uninformed agents to incorporate stale news into their forecasts, given the persistence of these shocks. This result offers a potential explanation for the empirical findings regarding stale information reported by Gilbert et al. (2012) and Tetlock (2011). The effect of this configuration on returns is a distribution with extremely low kurtosis and skewness, as shown in Table 4 and panel (c). The market is as efficient as possible, given the inability of uninformed agents to access private information. As a result, there is almost no room for extremely large returns or extreme losses.

We then analyze the second scale-free society, in which the most connected nodes are misinformed agents. The network in panel (a) of Fig. 8 has the same configuration as the previous one, with the only difference being the position of agents. Turning to the profit and accuracy analysis in panel (b), we can see that misinformation has successfully spread throughout the network. Misinformed agents now achieve positive and significantly high profits. This comes at the expense of uninformed agents, who are

[8] There are multiple example of investment platforms building on this feature. A famous one is the multi-asset investment platform eToro, in which users can directly copy the portfolio allocation of *top-performing investors*.

[9] It must be remarked that although until now we were using undirected networks, informed and misinformed agents were behaving in a dogmatic fashion. This is because a 0 prior variance in their beliefs implies non updating or posterior beliefs exactly equal to the prior.

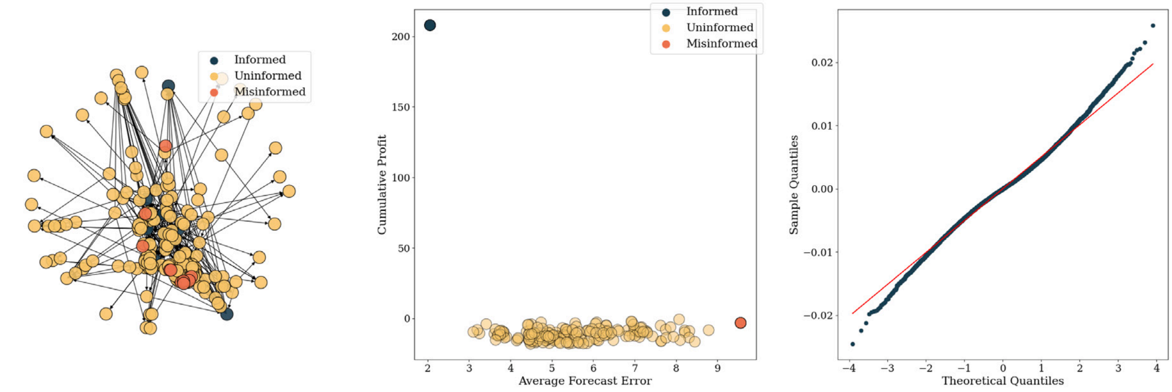

(a) Network Structure (b) Cumulative Profits (c) QQ Plot

Fig. 7. Barabasi–Albert graph, informed.
Panel (a): Scale Free Network structure. Each dot represents an agent, edges represents connections. Panel (b): average forecast error and profit for each agent. Panel (c): model based quantiles (blue) vs. theoretical (red) quantiles based on a Normal distribution with same mean and variance. Note: Data are obtained from 30 simulations with different stochastic seeds. **Network parameters** are: Alpha = 0.41, Beta = 0.54, Gamma = 0.05. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

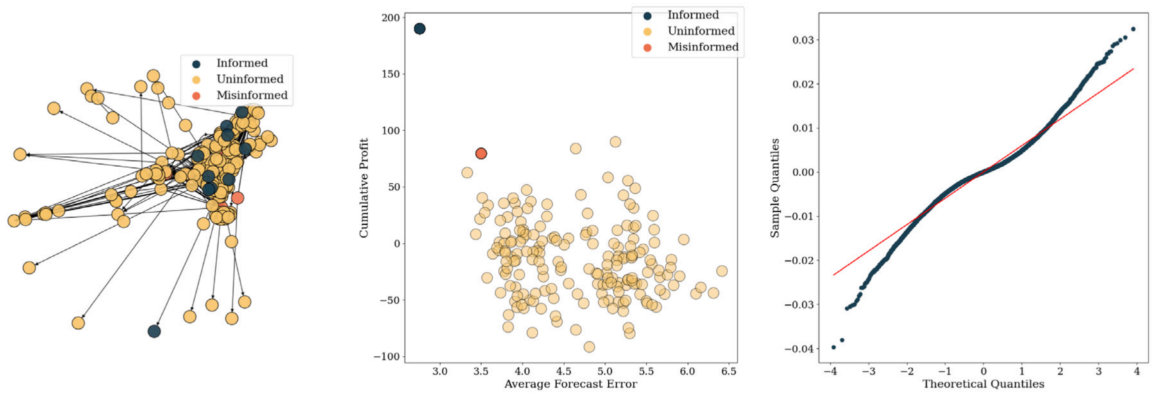

(a) Network Structure (b) Cumulative Profits (c) QQ Plot

Fig. 8. Barabasi–Albert graph, misinformed.
Panel (a): Scale Free Network structure. Each dot represents an agent, edges represents connections. Panel (b): average forecast error and profit for each agent. Panel (c): model based quantiles (blue) vs. theoretical (red) quantiles based on a Normal distribution with same mean and variance. Note: Data are obtained from 30 simulations with different stochastic seeds. **Network parameters** are: Alpha = 0.41, Beta = 0.54, Gamma = 0.05. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

both losing more on average and showing a high degree of inequality in their profits. Having, in most cases, connections only to misinformed individuals or other uninformed agents with stale misinformation makes their participation in the market extremely unfruitful. Panel (c) confirms that returns are again extremely leptokurtic, as reported in Table 4, and the skewness is now negative. This market is the least efficient of all the scenarios we have analyzed. However, some uninformed agents are still able to make profits. This is because the network structure allows certain agents to profit from future agents incorporating their beliefs, even though these beliefs are based on misinformation. The presence of social media platforms with scale-free structures, like Reddit or X, can lead to mechanisms that resemble pyramid or Ponzi schemes. In such schemes, even if participants realize they are involved, profits can still be made as long as new participants continue to join. Recent studies have shown that retail investors coordinating trades, not necessarily driven by fundamentals, have significantly impacted the prices and trading volumes of meme stocks, most notably GameStop, as documented in recent research (Costola et al., 2021; Hasso et al., 2022).

5. Conclusion

We have presented an Agent-Based Model of a financial market to explore how information and misinformation diffuse through social networks and shape asset prices and returns. In our framework, agents are embedded in a network and rely on their peers to

Table 4
Summary of moments for different scenarios.

	Profit of misinformed	Skewness	Kurtosis
Small World	−4.30	0.04	5.50
Stochastic Block Network	−3.17	0.08	7.85
Scale Free Informed	0.29	0.04	3.77
Scale Free Misinformed	85.58	−0.01	5.44

enhance their understanding of the underlying dividend process. We introduced a novel expectation formation mechanism, grounded in Bayesian updating, which allows agents to weigh multiple sources of information while departing minimally from the standard Full Information Rational Expectations (FIRE) benchmark. Through numerical simulations, we analyzed how different network structures affect market efficiency and asset return distributions.

The model replicates key empirical stylized facts observed in financial markets: returns display leptokurtosis, information is absorbed into prices with delays, and misinformation, while attenuated, still exerts influence. These features are consistent with documented patterns in the literature on return distributions and the gradual incorporation of news into prices (Huberman and Regev, 2001; Clarke et al., 2021). Moreover, we show that wealth inequality is amplified in networks where misinformed agents occupy central positions, highlighting how the structure of social influence can generate persistent disparities even in a zero-sum market.

While our analysis is positive in nature and does not endogenize network formation, it nonetheless yields important policy insights. First, our findings suggest that relying solely on past forecast accuracy as a filtering mechanism is not sufficient to prevent the spread of misinformation when misinformed agents have disproportionate communicative power. This has direct implications for both regulators and platform designers: fostering network topologies that emphasize peer-based connections, where information flows through trusted, reciprocal links, can help mitigate the distortive effects of misinformation. Social media platforms might consider mechanisms that dampen the influence of one-to-many broadcasting features and instead encourage more symmetric and verifiable interactions.

For retail investors, our results underscore the importance of being aware of the informational environment in which they are embedded. Understanding who influences their beliefs and how information circulates in their network can help them make better-informed decisions and avoid being swayed by prominent but inaccurate sources. Ultimately, our model provides a framework for thinking about how subtle frictions in information diffusion can give rise to observable market anomalies and unequal financial outcomes.

Code Availability Statement

All data used in the paper is publicly available. The code used to generate the results is available at https://github.com/danielTorren/Misinformation_financial_markets.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT in order to improve the readability and language of the manuscript. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Forward looking price

$$
\begin{aligned}
p_t &= R^{-1}\tilde{\mathbb{E}}_t\left(p_{t+1}+d_{t+1}\right) \\
&= R^{-1}\left[\tilde{\mathbb{E}}_t\left(p_{t+1}\right)+\tilde{\mathbb{E}}_t\left(d_{t+1}\right)\right] \\
&= R^{-1}\left[\tilde{\mathbb{E}}_t\left(R^{-1}\tilde{\mathbb{E}}_{t+1}\left(p_{t+2}+d_{t+2}\right)\right)+\tilde{\mathbb{E}}_t\left(d_{t+1}\right)\right] \\
&= R^{-1}\tilde{\mathbb{E}}_t(d_{t+1})+R^{-2}\tilde{\mathbb{E}}_t(d_{t+2})+R^{-2}\tilde{\mathbb{E}}_t(p_{t+2}) \\
&= R^{-1}\tilde{\mathbb{E}}_t(d_{t+1})+R^{-2}\tilde{\mathbb{E}}_t(d_{t+2})+R^{-2}\tilde{\mathbb{E}}_t(R^{-1}\mathbb{E}_{t+2}\left(p_{t+3}+d_{t+3}\right)) \\
&\vdots \\
&= \sum_{j=1}^{T} R^{-j}\tilde{\mathbb{E}}_t\left(d_{t+j}\right)+R^{-T}\tilde{\mathbb{E}}_t(p_{t+T}),
\end{aligned}
\tag{40}
$$

which implies, imposing $\lim_{T\to\infty} R^{-T}\tilde{\mathbb{E}}_t(p_{t+T}) = 0$ that in the limit $T \to \infty$

$$p_t = \sum_{j=1}^{\infty} R^{-j}\tilde{\mathbb{E}}_t\left(d_{t+j}\right). \tag{41}$$

Now we focus on the term $\tilde{\mathbb{E}}_t\left(d_{t+j}\right)$. We have that[10]

$$\tilde{\mathbb{E}}_t\left(d_{t+1}\right) = \tilde{\mathbb{E}}_t\left(d + \theta_{t+1} + \varepsilon_{t+1}\right) = d + \tilde{\mathbb{E}}_t\left(\theta_{t+1}\right) + \tilde{\mathbb{E}}_t\left(\varepsilon_{t+1}\right) = d + \tilde{\mathbb{E}}_t\left(\theta_{t+1}\right), \tag{42}$$

since $\tilde{\mathbb{E}}_t(\varepsilon_{t+1}) = 0$.

Similarly

$$\tilde{\mathbb{E}}_t\left(d_{t+2}\right) = \tilde{\mathbb{E}}_t\left(d + \theta_{t+2} + \varepsilon_{t+2}\right) = d + \tilde{\mathbb{E}}_t\left(\theta_{t+2}\right) + \tilde{\mathbb{E}}_t\left(\varepsilon_{t+2}\right) = d + \beta\tilde{\mathbb{E}}_t(\theta_{t+1}), \tag{43}$$

since $\tilde{\mathbb{E}}_t(\varepsilon_{t+2}) = 0$ and $\tilde{\mathbb{E}}_t(\theta_{t+2}) = \tilde{\mathbb{E}}_t(\beta\theta_{t+1} + \eta_{t+2})$, and in general

$$\tilde{\mathbb{E}}_t\left(d_{t+j}\right) = d + \beta^{j-1}\tilde{\mathbb{E}}_t(\theta_{t+1}) \quad \text{for all} \quad j \geq 1. \tag{44}$$

Therefore we have

$$p_t = \sum_{j=1}^{\infty} R^{-j}\left(d + \beta^{j-1}\tilde{\mathbb{E}}_t(\theta_{t+1})\right) = d\sum_{j=1}^{\infty} R^{-j} + \beta^{-1}\tilde{\mathbb{E}}_t(\theta_{t+1})\sum_{j=1}^{\infty}\left(\frac{\beta}{R}\right)^j. \tag{45}$$

Now we have

$$d\sum_{j=1}^{\infty} R^{-j} = d\sum_{j=0}^{\infty} R^{-j} - d = d\frac{R}{R-1} - d = \frac{d}{r}, \tag{46}$$

since $|R^{-1}| < 1$. Similarly

$$\begin{aligned}\beta^{-1}\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right)\sum_{j=1}^{\infty}\left(\frac{\beta}{R}\right)^j &= \beta^{-1}\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right)\sum_{j=0}^{\infty}\left(\frac{\beta}{R}\right)^j - \beta^{-1}\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right)\\ &= \beta^{-1}\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right)\frac{R}{R-\beta} - \beta^{-1}\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right) = \beta^{-1}\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right)\frac{\beta}{R-\beta}.\end{aligned} \tag{47}$$

So that finally

$$p_t = \frac{d}{r} + \frac{\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right)}{R-\beta}. \tag{48}$$

Appendix B. Relationship of updating to Kalman filter

Following the notation in chapter 13 of Hamilton (1994) we have the following state space representation for the observable component of dividends θ_t:

$$\begin{aligned}\text{(state equation)} \quad \xi_t &= F\xi_{t-1} + v_t,\\ \text{(measurement equation)} \quad y_t &= H\xi_t + w_t,\end{aligned} \tag{49}$$

in which all quantities are scalars, $\xi_t \equiv \theta_{t+1}$, $F \equiv \beta$, $H \equiv 1$. The variance–covariance matrix R associated with w_t is the scalar $\sigma^2_{j,t}$. In this contest the a priori variance covariance matrix is given by

$$P_{t|t-1} = \mathbb{E}(\xi_t - \hat{\xi}_{t|t-1})^2 = \sigma^2_\eta, \tag{50}$$

and the Kalman Gain

$$K_t = P_{t|t-1}H(H'P_{t|t-1}H + R)^{-1}, \tag{51}$$

collapses to

$$K_t = \frac{\sigma^2_\eta}{\sigma^2_\eta + \sigma^2_{j,t}}, \tag{52}$$

which is exactly the weight associated to the information received by source j in the case of being connected to source j only.

[10] Clearly if agents observe the realization of the stochastic component then $\tilde{\mathbb{E}}_t\left(\theta_{t+1}\right) = \theta_{t+1}$. We opt to keep the notation general because this allows us to simultaneously treat also agents that do not observe the realization of this noisy component.

Appendix C. Derivation of conditional variance

$$\tilde{\mathbb{V}}_t(p_{t+1} + d_{t+1}) = \tilde{\mathbb{V}}_t(d_{t+1}) + \tilde{\mathbb{V}}_t(p_{t+1}) + 2\tilde{\mathbb{COV}}_t(d_{t+1}, p_{t+1}). \tag{53}$$

The conditional variance of next period dividends is given by

$$\tilde{\mathbb{V}}_t(d_{t+1}) = \tilde{\mathbb{V}}_t(\theta_{t+1}) + \sigma_\varepsilon^2. \tag{54}$$

The conditional variance of next period price can be derived starting from the expression for p_{t+1} implied by Eq. (9)

We get

$$\tilde{\mathbb{V}}_t(p_{t+1}) = \tilde{\mathbb{V}}_t\left(\frac{d}{r} + \frac{\tilde{\mathbb{E}}_{t+1}(\theta_{t+2})}{R-\beta}\right) = \frac{\tilde{\mathbb{V}}_t(\tilde{\mathbb{E}}_{t+1}(\theta_{t+2}))}{(R-\beta)^2}. \tag{55}$$

In a similar fashion we can derive the expression for the covariance of the future price and dividend as

$$2\tilde{\mathbb{COV}}_t(d_{t+1}, p_{t+1}) = 2\tilde{\mathbb{COV}}_t\left(\theta_{t+1}, \frac{\tilde{\mathbb{E}}_{t+1}(\theta_{t+2})}{R-\beta}\right), \tag{56}$$

since d is a constant and ε_{t+1} is independent of any other variable. Rearranging we get Eq. (23).

Appendix D. Proofs

D.1. Proof of Proposition 1

The proof is by induction. First we prove the statement for the base case with only one information source. This collapses to the normal case of Bayesian updating with conjugate normal prior, therefore we have:

$$\mu_P = \frac{\mu_0\tau_0 + \mu_1\tau_1}{\tau_0 + \tau_1}, \tag{57}$$

$$\tau_P = \tau_0 + \tau_1. \tag{58}$$

Then we prove that if the statement holds for a generic K, then it holds also for $K+1$. If the statement holds for K, and the agents receives an extra source of information, the new posterior distribution has parameters:

$$\tau_P^{(K+1)} = \tau_P^{(K)} + \tau_{K+1} = \sum_{k=0}^{K+1} \tau_k,$$

and

$$\mu_P^{(K+1)} = \frac{\mu_P^{(K)}\tau_P^{(K)} + \mu_{K+1}\tau_{K+1}}{\tau_P^{(K)} + \tau_{K+1}} = \frac{\sum_{k=0}^{K+1}\left(\mu_k\tau_k\right)}{\tau_P^{(K+1)}}. \tag{59}$$

D.2. Proof of Proposition 2

Since $\sigma_P^2 = \frac{1}{\tau_P}$ it is enough to show that $\lim_{K\to\infty} \tau_P = +\infty$, and this follows from the definition $\tau_P = \sum_{k=0}^{K} \tau_k$ since $\tau_k > 0 \quad \forall k$.

Appendix E. Robustness to different stocks

While in the main text we focused on Tesla as a representative case, to assess robustness, we repeat the calibration exercise for all stocks in the S&P 500. As we can see in Fig. 9, the central finding that misinformation is associated with greater volatility than information, holds for 81% of companies in the index. Moreover, as illustrated in the figure, the estimated variances for Tesla are not outliers relative to the broader distribution, suggesting that the dynamics we study in the main body are not driven by idiosyncratic features of this particular asset.

Appendix F. Validation of the SNPE

To validate the model we generate synthetic data from the ABM and then we use the SNPE algorithm to recover the parameters. We fix $\sigma_\eta = \sigma_\nu = 0.7$ and keep the other parameters to their values in Table 1. We choose a uniform prior for the parameters with support $[0.1, 2]$ and we provide the result of 10 000 stochastic simulations to train the density estimator. The result is displayed in Fig. 10 and we can see that the posterior distribution is unimodal with median values 0.95 and 0.71 respectively for σ_η and σ_ν. The value of the first parameter is slightly off, but well within the confidence interval, and most likely due the fact that the posterior looks flat in the region of the true value.

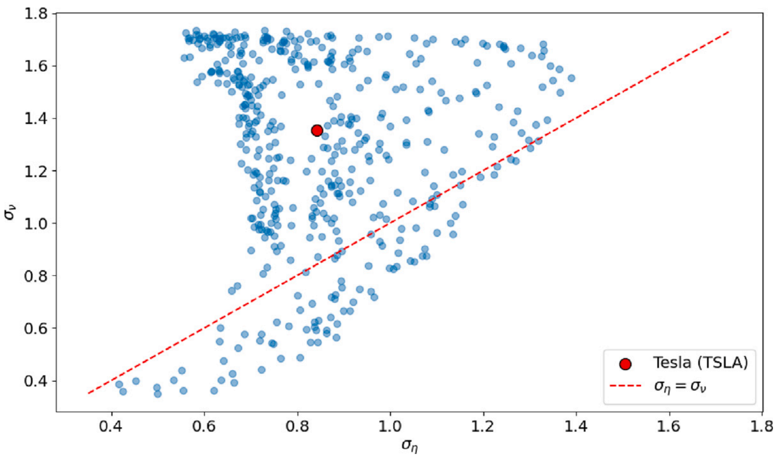

Fig. 9. Posterior median values for stocks in the S&P500.
Note: each dot represents posterior median values for σ_ν and σ_η for stocks in the S&P500 index. Above the red line the variance of the misinformation process is greater than that of the information process. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

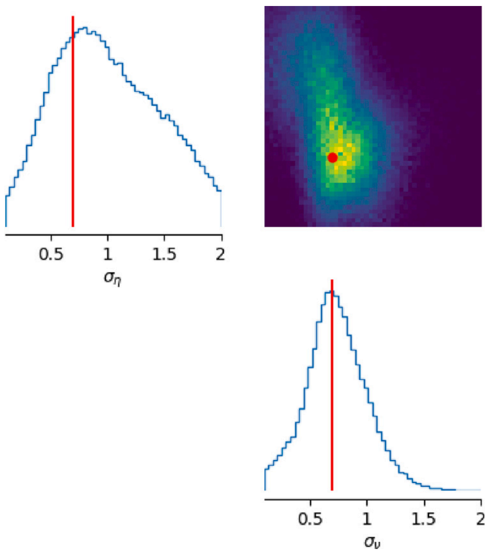

Fig. 10. Validation of the SNPE.
Note: The Figure shows the posterior distribution of the parameters obtained by the SNPE algorithm. The red lines represent the true values of the parameters. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Data availability

Link is shared in the paper.

References

Acemoglu, D., Ozdaglar, A., 2011. Opinion dynamics and learning in social networks. Dyn. Games Appl. 1 (1), 3–49.

Acemoglu, D., Ozdaglar, A., ParandehGheibi, A., 2010. Spread of (mis)information in social networks. Games Econom. Behav. 70 (2), 194–227.

Anufriev, M., Assenza, T., Hommes, C., Massaro, D., 2013. Interest rate rules and macroeconomic stability under heterogeneous expectations. Macroecon. Dyn. 17 (8), 1574–1604.

Anufriev, M., Gardini, L., Radi, D., 2020. Chaos, border collisions and stylized empirical facts in an asset pricing model with heterogeneous agents. Nonlinear Dynam. 102 (2), 993–1017.

Anufriev, M., Hommes, C., 2012. Evolutionary selection of individual expectations and aggregate outcomes in asset pricing experiments. Am. Econ. J.: Microecon. 4 (4), 35–64.

Assenza, T., Heemeijer, P., Hommes, C., Massaro, D., 2021. Managing self-organization of expectations through monetary policy: A macro experiment. J. Monet. Econ. 117, 170–186.

Aymanns, C., Foerster, J., Georg, C.-P., Weber, M., 2022. Fake News in Social Networks. Swiss Finance Institute Research Paper No. 22-58.

Banerjee, A.V., 1992. A simple model of herd behavior. Q. J. Econ. 107 (3), 797–817.

Barberis, N., Shleifer, A., Vishny, R., 1998. A model of investor sentiment. J. Financ. Econ. 49 (3), 307–343.

Benabou, R., Laroque, G., 1992. Using privileged information to manipulate markets: Insiders, gurus, and credibility. Q. J. Econ. 107 (3), 921–958.

Bertasiute, A., Massaro, D., Weber, M., 2020. The behavioral economics of currency unions: Economic integration and monetary policy. J. Econom. Dynam. Control 112, 103850.

Bertella, M.A., Silva, J.N., Correa, A.L., Sornette, D., 2021. The influence of confidence and social networks on an agent-based model of stock exchange. Complexity.

Biondo, A.E., 2020. Information versus imitation in a real-time agent-based model of financial markets. J. Econ. Interact. Coord. 15 (3), 613–631.

Black, F., 1986. Noise. J. Financ. 41 (3), 528–543.

Blankespoor, E., deHaan, E., Marinovic, I., 2020. Disclosure processing costs, investors' information choice, and equity market outcomes: A review. J. Account. Econ. 70 (2–3), 101344.

Board, S., Meyer-ter Vehn, M., 2021. Learning dynamics in social networks. Econometrica 89 (6), 2601–2635.

Bollobás, B., Borgs, C., Chayes, J., Riordan, O., 2003. Directed scale-free graphs. In: Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms. SODA '03, Society for Industrial and Applied Mathematics, USA, pp. 132–139.

Brock, W.A., Hommes, C.H., 1997. A rational route to randomness. Econometrica 65 (5), 1059–1095.

Brock, W.A., Hommes, C.H., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. J. Econom. Dynam. Control 22 (8–9), 1235–1274.

Buechel, B., Hellmann, T., Klößner, S., 2015. Opinion dynamics and wisdom under conformity. J. Econom. Dynam. Control 52, 240–257.

Chetty, R., 2006. A new method of estimating risk aversion. Am. Econ. Rev. 96 (5), 1821–1834.

Chiarella, C., 1992. The dynamics of speculative behaviour. Ann. Oper. Res. 37 (1), 101–123.

Clarke, J., Chen, H., Du, D., Hu, Y.J., 2021. Fake news, investor attention, and market reaction. Inf. Syst. Res. 32 (1), 35–52.

Collin-Dufresne, P., Fos, V., 2016. Insider trading, stochastic liquidity, and equilibrium prices. Econometrica 84 (4), 1441–1475.

Cont, R., Bouchaud, J.-P., 2000. Herd behavior and aggregate fluctuations in financial markets. Macroecon. Dyn. 4 (2), 170–196.

Costola, M., Iacopini, M., Santagiustina, C.R., 2021. On the "mementum" of meme stocks. Econom. Lett. 207, 110021.

Dasaratha, K., Golub, B., Hak, N., 2022. Learning from neighbours about a changing state. Rev. Econ. Stud. 90 (5), 2326–2369.

Day, R.H., Huang, W., 1990. Bulls, bears and market sheep. J. Econ. Behav. Organ. 14 (3), 299–329.

De Grauwe, P., Grimaldi, M., 2006. The Exchange Rate in a Behavioral Finance Framework. Princeton University Press.

De Long, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990. Noise trader risk in financial markets. J. Political Econ. 98 (4), 703–738.

Degroot, M.H., 1974. Reaching a consensus. J. Amer. Statist. Assoc. 69 (345), 118–121.

Dellavigna, S., Pollet, J.M., 2009. Investor inattention and friday earnings announcements. J. Financ. 64 (2), 709–749.

DeMarzo, P.M., Vayanos, D., Zwiebel, J., 2003. Persuasion bias, social influence, and unidimensional opinions. Q. J. Econ. 118 (3), 909–968.

Di Francesco, T., Hommes, C., 2025. Sentiment-driven speculation in financial markets with heterogeneous beliefs: A machine learning approach. J. Econom. Dynam. Control 175, 105092.

Diks, C., Dindo, P., 2008. Informational differences and learning in an asset market with boundedly rational agents. J. Econom. Dynam. Control 32 (5), 1432–1465.

Dyer, J., Cannon, P., Farmer, J.D., Schmon, S.M., 2024. Black-box Bayesian inference for agent-based models. J. Econom. Dynam. Control 161, 104827.

Franke, R., Westerhoff, F., 2012. Structural stochastic volatility in asset pricing dynamics: Estimation and model contest. J. Econom. Dynam. Control 36 (8), 1193–1211.

Franke, R., Westerhoff, F., 2016. Why a simple herding model may generate the stylized facts of daily returns: explanation and estimation. J. Econ. Interact. Coord. 11 (1), 1–34.

Gale, D., Kariv, S., 2003. Bayesian learning in social networks. Games Econom. Behav. 45 (2), 329–346.

Gardini, L., Radi, D., Schmitt, N., Sushko, I., Westerhoff, F., 2022. Causes of fragile stock market stability. J. Econ. Behav. Organ. 200, 483–498.

Gardini, L., Radi, D., Schmitt, N., Sushko, I., Westerhoff, F., 2025. On boom-bust stock market dynamics, animal spirits, and the destabilizing nature of temporarily attracting virtual fixed points. Macroecon. Dyn. 29, e35.

Gerotto, L., Pellizzari, P., Tolotti, M., 2019. Asymmetric information and learning by imitation in agent-based financial markets. In: De La Prieta, F., González-Briones, A., Pawleski, P., Calvaresi, D., Del Val, E., Lopes, F., Julian, V., Osaba, E., Sánchez-Iborra, R. (Eds.), Highlights of Practical Applications of Survivable Agents and Multi-Agent Systems. the PAAMS Collection. Springer International Publishing, Cham, pp. 164–175.

Gilbert, T., Kogan, S., Lochstoer, L., Ozyildirim, A., 2012. Investor inattention and the market impact of summary statistics. Manag. Sci. 58 (2), 336–350.

Golub, B., Jackson, M.O., 2010. Naive learning in social networks and the wisdom of crowds. Am. Econ. J.: Microecon. 2 (1), 112–149.

Grossman, S.J., Stiglitz, J.E., 1980. American economic association on the impossibility of informationally efficient markets. Am. Econ. Rev. 70 (3), 393–408.

Hamilton, J.D., 1994. Time Series Analysis. Princeton University Press.

Hasso, T., Müller, D., Pelster, M., Warkulat, S., 2022. Who participated in the GameStop frenzy? Evidence from brokerage accounts. Financ. Res. Lett. 45, 102140.

Heemeijer, P., Hommes, C., Sonnemans, J., Tuinstra, J., 2009. Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. J. Econom. Dynam. Control 33 (5), 1052–1072.

Hirshleifer, D., Lim, S.S., Teoh, S.H., 2009. Driven to distraction: Extraneous events and underreaction to earnings news. J. Financ. 64 (5), 2289–2325.

Holland, P.W., Laskey, K.B., Leinhardt, S., 1983. Stochastic blockmodels: First steps. Soc. Netw. 5 (2), 109–137.

Hommes, C., Massaro, D., Weber, M., 2019. Monetary policy under behavioral expectations: Theory and experiment. Eur. Econ. Rev. 118, 193–212.

Huberman, G., Regev, T., 2001. Contagious speculation and a cure for cancer: A nonevent that made stock prices soar. J. Financ. 56 (1), 387–396.

Iori, G., 2002. A microsimulation of traders activity in the stock market: The role of heterogeneity, agents' interactions and trade frictions. J. Econ. Behav. Organ. 49 (2), 269–285.

Jegadeesh, N., Titman, S., 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. J. Financ. 48, 65–91.

Kanoria, Y., Tamuz, O., 2013. Tractable bayesian social learning on trees. IEEE J. Sel. Areas Commun. 31 (4), 2721–2725.

Khashanah, K., Alsulaiman, T., 2016. Network theory and behavioral finance in a heterogeneous market environment. Complexity 21.

Kopányi-Peuker, A., Weber, M., 2020. Experience does not eliminate bubbles: Experimental evidence. Rev. Financ. Stud. 34 (9), 4450–4485.

Kyle, A.S., 1985. Continuous auctions and insider trading. Econometrica 53 (6), 1315–1335.

Lux, T., 1998. The socio-economic dynamics of speculative markets: Interacting agents, chaos, and the fat tails of return distributions. J. Econ. Behav. Organ. 33 (2), 143–165.

Massaro, D., 2013. Heterogeneous expectations in monetary DSGE models. J. Econom. Dynam. Control 37 (3), 680–692.

Mueller-Frank, M., Neri, C., 2013. Social learning in networks: Theory and experiments. Work. Pap..

Orléan, A., 1995. Bayesian interactions and collective dynamics of opinion: Herd behavior and mimetic contagion. J. Econ. Behav. Organ. 28 (2), 257–274.

Panchenko, V., Gerasymchuk, S., Pavlov, O.V., 2013. Asset price dynamics with heterogeneous beliefs and local network interactions. J. Econom. Dynam. Control 37 (12), 2623–2642.

Papamakarios, G., Murray, I., 2016. Fast epsilon-free inference of simulation models with Bayesian conditional density estimation. In: Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., Garnett, R. (Eds.), Advances in Neural Information Processing Systems. Vol. 29, Curran Associates, Inc..

Peng, L., Xiong, W., 2006. Investor attention, overconfidence and category learning. J. Financ. Econ. 80 (3), 563–602.
Rusinowska, A., Taalaibekova, A., 2019. Opinion formation and targeting when persuaders have extreme and centrist opinions. J. Math. Econom. 84, 9–27.
Schmitt, N., Westerhoff, F., 2017. Heterogeneity, spontaneous coordination and extreme events within large-scale and small-scale agent-based financial market models. J. Evol. Econ. 27 (5), 1041–1070.
Sobol, I.M., 2001. Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. Math. Comput. Simulation 55 (1–3), 271–280.
Tedeschi, G., Iori, G., Gallegati, M., 2012. Herding effects in order driven markets: The rise and fall of gurus. J. Econ. Behav. Organ. 81 (1), 82–96.
Tetlock, P.C., 2011. All the news that's fit to reprint: Do investors react to stale information? Rev. Financ. Stud. 24 (5), 1481–1512.
Tramontana, F., Westerhoff, F., Gardini, L., 2010. On the complicated price dynamics of a simple one-dimensional discontinuous financial market model with heterogeneous interacting traders. J. Econ. Behav. Organ. 74 (3), 187–205.
Tramontana, F., Westerhoff, F., Gardini, L., 2013. The bull and bear market model of Huang and day: Some extensions and new results. J. Econom. Dynam. Control 37 (11), 2351–2370.
Vozlyublennaia, N., 2014. Investor attention, index performance, and return predictability. J. Bank. Financ. 41 (1), 17–35.
Watts, D.J., Strogatz, S.H., 1998. Collective dynamics of 'small-world' networks. Nature 393 (6684), 440–442.
Wu, S., He, J., Li, S., 2018. Effects of fundamentals acquisition and strategy switch on stock price dynamics. Phys. A 491, 799–809.