

RESEARCH ARTICLE

Herding or wisdom of the crowd? Controlling efficiency in a partially rational financial market

Fabio Della Rossa^{1,2}, Lorenzo Giannini², Pietro DeLellis²*

1 Department of Electronic, Information, and Bioengineering, Politecnico di Milano, Milan, Italy, **2** Department of Electrical Engineering and Information Technology, University of Naples Federico II, Naples, Italy

✉ These authors contributed equally to this work.

* pietro.delellis@unina.it



Abstract

Herding has often been blamed as one of the possible causes of market instabilities, ultimately yielding to bubbles and crashes. On the other hand, researchers hypothesized that financial systems may benefit from the so-called *wisdom of the crowd*. To solve this apparent dichotomy, we leverage a novel financial market model, where the agents form their expectations by combining their individual return estimation with the expectations of their neighbors. By establishing a link between herding, sociality, and market instabilities, we point out that the emergence of collective decisions in the market is not necessarily detrimental. Indeed, when all the agents tend to conform their expectations to those of one or few leaders, herding might dramatically reduce market efficiency. However, when each agent accounts for a plurality of opinions, thus following the wisdom of the crowd, market dynamics become efficient. Following these observations, we propose two alternative control strategies to reduce market instability and enhance its efficiency.

OPEN ACCESS

Citation: Della Rossa F, Giannini L, DeLellis P (2020) Herding or wisdom of the crowd? Controlling efficiency in a partially rational financial market. PLoS ONE 15(9): e0239132. <https://doi.org/10.1371/journal.pone.0239132>

Editor: Hocine Cherifi, University of Burgundy, FRANCE

Received: March 2, 2020

Accepted: August 27, 2020

Published: September 11, 2020

Copyright: © 2020 Della Rossa et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: All relevant data are within the manuscript.

Funding: We wish to thank the University of Naples Federico II and the Compagnia di San Paolo, Istituto Banco di Napoli - Fondazione for supporting this research under the grant STAR 2018, project ACROSS. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

Introduction

In neoclassical economics, the efficient market assumption prescribes that the market price of an asset fully reflects all available information [1]. This implies that financial agents form rational expectations about future price variations, thus driving the market price towards the correct *fundamental value* of the traded asset [2]. This theory has been severely criticized in the last decades, since it failed to explain periods when the market consistently misestimates the value of an asset [3]. A possible explanation of this empirical evidence is that the decision-making in humans is not perfectly rational, and might be subservient to emotions and cognitive biases [4, 5]. Also, individual decisions are not taken independent of the opinions and behaviors of the other investors, see for instance the behavior of trend followers, which might contribute to trigger market instabilities [6]. Several financial bubbles are induced by similar mechanisms, as in the so-called *dot-com boom* that took place at the beginning of the 21st century, where investors overlooked traditional metrics of profitability [7], leading to a rapid increase and subsequent decrease of the share prices of the IT companies.

Considering their potentially disruptive socio-economic impact, identifying the mechanisms underlying the emergence of *market bubbles* is paramount [8, 9]. Since market crashes are observed in the presence of a significant difference between the market prices and the real value of the traded assets, their emergence is not compatible with the traditional efficient market assumption. A crucial step towards understanding the market dynamics has been made by behavioral economics, which contributed a better characterization of the human decision process, pervaded by irrationality and uncertainties [10]. This yielded a more accurate atomized characterization of the human behavior that, paired with the increased computational power, allowed for the development of *artificial financial markets* trying to reproduce in a simulated environment the features of real financial markets [11]. A considerable research effort has been devoted to decipher the effect of imitative mechanisms in financial markets, see e.g. [12–14]. Indeed, the financial agents, not being aware of the of the intrinsic value of a resource, may leverage the opinion of her peers to form their expectation. In the last decade, the presence of imitative behaviors in trading has been magnified by social trading platforms like *eToro* and *Zulutrade* [15], which allow investors to observe the trading patterns of their peers and copy their investment strategies.

In the literature, there is no consensus on the impact of imitation on the market dynamics. Often, when the emergence of a collective decision in the market is identified as one of the possible causes of market instabilities, the derogatory word *herding* is used [16, 17]. The use of this term to indicate the tendency to (blindly) imitate the behavior of other individuals is rather consistent in the literature [18] and, interestingly, has a negative connotation also in other research fields, including in the studies on crowd formation during emergencies [19, 20]. In financial markets, herding has been viewed as a cascading imitation process that reinforces agent expectations [21], thus causing an increment in volatility [22]. On the other hand, considering the inherent uncertainties of financial markets, imitating the behavior of the average investor, with the agents following the so-called *wisdom of the crowd*, an ensemble of diversely informed individuals may outperform experts, and contribute to steer the asset price towards its fundamental value, thus enhancing market efficiency. Examples of this phenomenon can be found in prediction markets [23], see e.g. [24, 25]. A first attempt to discriminate between a detrimental herding and a *rational adaptation* has been made in [26], where the complexity of the agent individual behavior was neglected to focus instead on the dynamic evolution of the network structure.

This work leverages the tools of complex network theory to explain the apparent dichotomy between herding and crowd wisdom. In particular, we employ a recent financial market model [27] to elucidate the relationship between sociality, imitation, and market efficiency. Specifically, we consider a centralized market with a double auction order book, and model the social interactions among the agents through a graph. The agents do not possess a perfect knowledge of the current fundamental price, but they form their own prediction, which is then adjusted based on the predictions of the neighboring agents. By varying the degree of sociality among the agents, we observe that, when all the traders conform their expectations to those of a small subset of investors, then herding is associated to a decrease of market efficiency. On the contrary, in the absence of a dominant opinion in the market, its efficiency may benefit from the social interaction among the investors. Driven by these numerical observations, we devised two alternative control strategy, which can be readily translated into market policies in social trading platforms, where the interaction topology among the agents can be manipulated. The first control strategy consists in isolating (i.e. disconnecting) the most influential (i.e. the wealthiest) traders by not disclosing their expectations, while the second is a *data anonymization* strategy, where the expectations are disclosed to neighboring agents, but not their identities. A thorough statistical analysis illustrates how both strategies are effective in enhancing

market efficiency even in the presence of a strong social influence among the investors. Indeed, we observe an instance of crowd wisdom, as the trading patterns become strongly correlated, but the market price tend to closely follow the fundamental price, this reflecting in a healthy and efficient market.

Methods

Double auction markets

In double auction markets [28], a centralized entity, i.e., the market institution, collects all the bids entered by the buyers and the offers submitted by the sellers, and store them in an *order book*. The bids (offers) of the traders consist in their willingness to buy (sell) a specific amount of goods at a *limit price*, which is the highest (lowest) price they are willing to accept for the trade. The order is stored in the book until another trader is willing to match the same price request, and it is *executed*. An order that is still not executed is called *limit order*. When a trader is willing to buy (sell) at a price another agent has already submitted a sell (buy) order, then her order is called *market order*, and is immediately executed. The highest of all the buy limit orders is the *bid price* p_{bid} , while the lowest sell limit order is the *ask price* p_{ask} . Market orders determine the execution of limit orders with limit price equal to the current bid and ask prices. The price of the last transaction is the current asset price p_{cur} , while the asset price at the beginning of the k -th day of trading is called the *market price* $p(k)$.

Artificial financial market model

Here, we describe in details a recent agent-based model of a single-asset double auction market, which was first presented in [27]. The market is populated by a set of N financial agents, who submit market or limit orders depending on their expectation about the future prices of the asset. The wealth $W_i(k)$ of agent i at day k depends on the quantity $a_i(k)$ of the asset that she possesses, and her current cash availability $c_i(k)$. Indeed,

$$W_i(k) = c_i(k) + a_i(k)p(k), \quad (1)$$

for $i = 1, \dots, N, k \in \mathbb{Z}^+$. Note that, when evaluating the wealth of an agent, we neglect price impact (i.e. the correlation between an incoming order and the subsequent price change [29]), and only consider the market price $p(k)$ at the beginning of the k -th trading day.

A key feature of this model is that it explicitly accounts for the social interaction among the agents. As in social trading platforms, our market model assumes that traders are aware of the behavior of a subset of their peers. Specifically, the agents are supposed to interact on a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} identifies the set of financial agents, while the presence of an edge $(i, j) \in \mathcal{E}$ implies that agent i has an influence over agent j . Indeed, when forming her expectations, the i -th agent only takes into account the decisions of the agents in her *in-neighborhood*

$$\mathcal{N}_i^{\text{in}} = \{i \in \mathcal{V} : (j, i) \in \mathcal{E}\}.$$

Every day, the sequence of trading is randomly extracted, and each agent performs a prediction of the future returns and then decides to submit a market or limit order in the book.

Return prediction. At the k -th day of trading, the i -th agent makes a prediction $\hat{r}_i(k)$ of the expected log-return of the traded asset, obtained as a convex combination of the agent's individual prediction $\hat{r}_{f,i}(t)$ and a social prediction $\hat{r}_{s,i}(t)$, that is a weighted average of the predictions of her in-neighbors.

Individual prediction. Different from the other market variables, the price prediction performed by each agent continuously updates with time. Namely, the dynamics of the prediction $\hat{p}_{f,i}$ of the fundamental price p_f performed by the i -th agent is modelled by the following Ornstein-Uhlenbeck process:

$$\partial \hat{p}_{f,i}(t) = \gamma_i(p_f(t) - \hat{p}_{f,i}(t))\partial t + \sigma_i \partial B_i(t), \tag{2}$$

where $t \in \mathbb{R}^+$, $\gamma_i(p_f(t) - \hat{p}_{f,i}(t))\partial t$ is a mean reversion term that captures the tendency towards a correct estimation of the fundamental price, $\partial B_i(t) \sim \mathcal{N}(0, 1)$ is a Wiener process modeling the inherent uncertainty affecting the estimation; the parameter γ_i determines the intensity of attraction towards p_f , which can be considered as a measure of the agent *expertise*, that is, her ability in estimating the fundamental price. On the other hand, the higher is the *insecurity* of the agent, the higher the variance σ_i^2 of the Wiener process will be.

The goal of model (2) is to capture the tradeoff between the tendency of a financial agent i) to accurately identifying the fundamental price (this is captured by the mean reverting term of Eq (2)) and ii) to deviate from it (captured by the Wiener process). The tradeoff between these two terms is regulated by the ratio between the agent expertise γ_i and her insecurity σ_i . Note that, to simulate the artificial financial market, we then need to assume a given fundamental value $p_f(t)$. In agreement with past works [30, 31], we take it as a realization of a geometric Wiener process. However, this is not assumed to be known by the agent.

In our study, we assume that all the prices and their estimations are positive. This is enforced through our model implementation (see the subsection Numerical Setup), where we take the fundamental price $p_f(t)$ at any time instant t sufficiently large if compared with the variance of the Wiener process, such that in all our repetitions $\hat{p}_f(t)$ is also positive. Therefore, during the k -th day of trading, the expected log-return is then estimated as

$$\hat{r}_{f,i}(k) = \log\left(\frac{\hat{p}_{f,i}(k)}{p_{cur}^i(k)}\right), \tag{3}$$

where $p_{cur}^i(k)$ is the asset price before agent i participates to the k -th trading session.

Social prediction. As in social trading, we assume that the agents, when investing, take into account the behavior of a subset of their peers. In our model, the behavior of an agent is captured by her expected log-return, which influences both price and demand determination. Therefore, if the set of her in-neighbors \mathcal{N}_i^{in} is non-empty, agent i , when deciding how to invest (i.e. computing her expected return), performs a weighted average of the expected log-returns of her in-neighbors, thus building her *social prediction* on the returns. The weights are selected so that the most successful traders, which in our model are identified as the wealthiest, are weighted the most. In real markets, different proxies of success can be considered. Namely,

$$\hat{r}_{s,i}(k) = \frac{1}{W_{in}(k)} \left(\sum_{j \in \mathcal{N}_i^{in}} W_j(k) \hat{r}_j(k-1) \right), \tag{4}$$

where $W_{in}(k)$ is the sum of the wealth of the in-neighbors of i at the beginning of day k . Note that, since we assumed $p(k) > 0$, Eq (1) ensures that $W_j(k)$ for all $j \in \mathcal{V}$ and $k \in \mathbb{Z}_{\geq 0}$, in turn guaranteeing that Eq (4) is well-posed.

Overall prediction. Agent i then computes the future log-return of the traded asset as a convex combination of the individual and social predictions:

$$\hat{r}_i(k) = (1 - \beta_i(k))\hat{r}_{f,i}(k) + \beta_i(k)\hat{r}_{s,i}(k), \tag{5}$$

where $\beta_i(k)$ is zero when $\mathcal{N}_i^{\text{in}} = \emptyset$, otherwise it is computed as

$$\beta_i(k) = \zeta \frac{W_{\text{in}}(k)}{W_{\text{in}}(k) + W_i(k)}.$$

The constant parameter $\zeta \in [0, 1]$ is the network *sociability*, which balances the relevance of $\hat{r}_{s,i}$ against the individual prediction of the return $\hat{r}_{f,i}$. The factor $W_{\text{in}}(k)/(W_{\text{in}}(k) + W_i(k))$ accounts for the fact that wealthier (i.e. more successful) agents will be less prone to consider the opinion of their neighbors. Note that $\zeta = 0$ and $\zeta = 1$ are representative of the cases of stubborn non-social agents, which neglect the opinions of the others, and of insecure social agents, who tend to disregard their individual predictions in favor of those of their neighbors, respectively.

Placing an order. *Determining the theoretical demand.* Following the approach presented in [30], agent i decides the quantity to buy/sell by maximizing the the expected value of her utility function, that is, the theoretical demand π_i of agent i at time k is given by

$$\pi(k) = \arg \max_{\pi} \mathbb{E}[U^i(W_i(k, \pi))], \tag{6}$$

where $W_i(k, \pi) = c(k) + \pi p(k)$. Two of the main classes of utility functions typically employed to model the agent behaviors are the CARA (constant absolute risk aversion) and CRRA (constant relative risk aversion) functions [32], which are defined as

$$U_{\text{CARA}}^i(W_i(k, \pi), \alpha_i) = -e^{-\alpha_i W_i(k, \pi)} \tag{7}$$

and

$$U_{\text{CRRA}}^i(W_i(k, \pi), \varsigma_i) = (W_i(k, \pi)^{\varsigma_i} - 1)/\varsigma_i, \tag{8}$$

respectively; the quantities α_i and $1 - \varsigma_i$ represent the risk aversion of the i -th agent. In the absence of dividends and risk-free assets, an approximate solution of (6) is

$$\pi_i(k) = \frac{\hat{r}_i(k)}{\varrho_i \hat{V}_i(k) p_{\text{cur}}^i(k)}, \tag{9}$$

where the risk aversion ϱ_i is equal to α_i if a CARA utility function is selected, while it is $1 - \varsigma_i$ if a CRRA utility function is used instead, see the Appendixes of [30] and [33] for the formal derivations; $\hat{V}_i(k)$ is the estimated volatility at the k -th day of trading, computed from the returns of the asset in the past τ_i trading sessions, that is,

$$\hat{V}_i(k) = \frac{1}{\tau_i} \sum_{\delta=1}^{\tau_i} (r(k - \delta) - \bar{r}_i(k))^2, \tag{10}$$

where $r(k) = \ln(p(k)/p(k - 1))$ is the spot return of the asset price at day $k - i$, while $\bar{r}_i(k)$ is the average spot return in the trading sessions $k - 1, \dots, k - \tau_i$. The theoretical demand $\pi_i(k)$ is then compared with the quantity $a_i(k)$ that the agent currently possesses to determine whether a buy or sell order has to be submitted. Note that, as in [30, 33], an agent can neither leverage debt to buy more shares, nor short sale, that is, she can only sell shares of the asset she currently owns.

Price determination. Similar to [34], the price set by the i -th agent is distributed around the bid (in case of a buy order) or ask price (in case of a sell order). Namely,

$$p_i(k) = \begin{cases} p_{\text{bid}}^i(k) - (\eta_i(k) - m), & \text{if } \pi_i(k) > a_i(k) \quad (\text{buy order}), \\ p_{\text{ask}}^i(k) + (\eta_i(k) - m), & \text{if } \pi_i(k) < a_i(k) \quad (\text{sell order}), \end{cases} \tag{11}$$

where $p_{bid}^i(k)$ and $p_{ask}^i(k)$ are the bid and ask price before agent i participates to the k -th trading session, and $\eta_i(k)$ is the realization of a log-normal random variable with meta-parameters defined as in [34], that is, with mean 7 and standard deviation 10; the shift parameter m is the median of the distribution, and regulates the tradeoff between market and limit orders. The log-normal distribution is selected in accordance with the empirical observations on the order distribution in the book: counter-intuitively, submitted orders are not always market orders or limit orders slightly above the best quote, but may also strongly deviate from p_{bid}^i (or p_{ask}^i) [35]. This behaviour has been explained in line with the trading strategies adopted by some traders, as for instance the *stop loss*. The strategy consists in submitting an order in the opposite direction of the current position to cap the maximum possible loss in case of a sudden adverse price variation.

Determining the order size. Agent i then compares the theoretical demand $\pi_i(k)$ with the current shares $a_i(k)$ she possesses, and with her current resource availability $c_i(k)$, thus deriving the order size $d_i(k)$ as

$$d_i(k) = \begin{cases} \lfloor \min(\pi_i(k) - a_i(k), c_i(k)/p_i(k)) \rfloor & \text{if } \pi_i(k) > a_i(k) \quad (\text{buy order}), \\ \lfloor \min(a_i(k) - \pi_i(k), a_i(k)) \rfloor & \text{if } \pi_i(k) < a_i(k) \quad (\text{sell order}), \end{cases} \tag{12}$$

where $\lfloor \cdot \rfloor$ is the floor function. Subsequently, agent i places a buy (or sell) order of $d_i(k)$ units of the asset at price $p_i(k)$, which is registered in the book.

Measuring herding and efficiency

To adequately quantify the intuitive concepts of herding and market efficiency, we introduce the following metrics:

- denoting $\hat{p}_i(k) = p_{cur}^i(k) \exp(\hat{r}_i(k))$ the price estimation performed by agent i , and $[t_0, T]$ the window of trading under analysis, we define the percent **herding intensity** H as

$$H = 100 \frac{\|\rho\|_2 - 1}{N - 1},$$

where ρ is the sample correlation matrix of the time series $\hat{p}_1(t_0), \dots, \hat{p}_i(T)$ of all agents $i = 1, \dots, N$, whose element lm is

$$\rho_{lm} = \frac{\sigma_{lm}^2}{\sigma_{ll}\sigma_{mm}},$$

with

$$\sigma_{lm} = \frac{1}{T - 1} \left(\sum_{k=t_0}^T (\hat{p}_l(k) - \langle \hat{p}_l \rangle) (\hat{p}_m(k) - \langle \hat{p}_m \rangle) \right)^{1/2},$$

and $\langle \hat{p}_i \rangle$ being the average \hat{p}_i in the trading window under analysis. Note that $0 \leq H \leq 100$, with $H = 0$ in the absence of correlation, and $H = 100$ when all the agents perform the same estimation throughout the trading window, that is, when $\hat{p}_i(k) = \hat{p}_j(k)$ for all (i, j) and for all $k \in [t_0, T]$.

- introducing the error $e(k) := p(k) - p_f(k)$ between the market and fundamental prices, we define the **market inefficiency** \bar{e} as

$$\bar{e} = \frac{1}{T} \sqrt{\sum_{t=t_0}^T e(k)^2}.$$

Results

Here, we numerically investigate the properties of the model presented above to elucidate the relationship between sociability, herding, and market efficiency. The outcome of this analysis is then leveraged to propose market policies that enhance its efficiency. Finally, we evaluate how variations in the maximum risk aversion of the traders impact on market dynamics.

Numerical setup

We consider an artificial financial market composed of $N = 100$ agents, whose pattern of interactions is captured by a Barabasi-Albert scale-free network [36] with average degree 4. A scale-free network topology has been selected as it is typically observed both in financial networks and in other social or biological contexts [37–40]. In agreement with past works, the fundamental price $p_f(t)$ is taken as a realization of a geometric Wiener process [30, 31]. The agents’ initial capital, wealth and return expectation, as well as the random distributions from which the model parameters are extracted, are reported in Table 1, and each simulation considers a time frame of $T = 1000$ trading sessions. Since the order book is empty at the beginning of each simulation, we consider $p_{cur}^i(0) = p_{bid}^i(0) = p_{ask}^i(0) = p_f(0)$, and we start computing herding intensity and market efficiency from $t_0 = \lfloor 3T/10 \rfloor$. In any trading session in which no buy (sell) orders are present in the book, both p_{bid}^i (p_{ask}^i) and p_{cur}^i are set to p_{ask}^i (p_{bid}^i). As wealth is considered as a proxy of success, the parameters have been set so that the wealthiest node is also the most influential, that is, the node, say i , with the highest out-degree d_i^{out} , defined as the the cardinality of the out-neighbors set

$$\mathcal{N}_i^{out} = \{i \in \mathcal{V} : (i, j) \in \mathcal{E}\}.$$

To test for significant variations of the metrics H and \bar{e} depending on the level of sociality or on the adopted market policies, one-way ANOVAs are employed in the following analyses. In particular, when the data in the group are non-normal (normality is verified with a Lilliefors test [41]) or when the two groups are heteroskedastic (different variance between the two groups, verified with the Bartlett test [42]), suitable transformations of the data have been applied before performing the ANOVA test.

Table 1. Market initialization and selected distribution for each of the parameters.

Variable	Initial value	Parameter	Distribution
$c_i(k)$	$3000(1 + d_i^{out})^2$	ϱ_i	$\mathcal{U}(0, 0.005)$
$a_i(k)$	100	γ_i	$\mathcal{N}(0.5, 0.1^2)$
$\hat{r}_i(k)$	0	τ_i	$\mathcal{N}(50, 10^2)$
		σ_i	$\mathcal{N}(5, 1)$

<https://doi.org/10.1371/journal.pone.0239132.t001>

Sociability, herding, and market behavior

To investigate how the mutual influence among the investors affects the market dynamics, we vary the sociability ζ in the set $[0, 1]$ with step 0.1 and, for each value of ζ , run 100 simulations to assess the statistical relevance of the observed results. For each value of $\zeta \neq 0$, we performed a one-way ANOVA to test whether herding intensity and efficiency were different from the case $\zeta = 0$, where sociability is absent, and we always obtained statistical significance, with p -values smaller than 0.005 (except for herding when $\zeta \leq 0.2$). In Fig 1, we report the box-plots of H and \bar{e} as a function of the sociability parameter ζ . We notice that, while the herding intensity monotonically increases with the sociability, its impact on \bar{e} is less trivial. Indeed, an increment in ζ is initially beneficial for the efficiency, while when $\zeta \geq 0.8$ market efficiency dramatically drops.

To further illustrate the relationship between herding and efficiency, we report in Fig 2 the time series of the expected prices \hat{p}_i of three sample simulations of our model with $\zeta = 0, 0.5$, and 1, respectively. In the absence of sociability (Fig 2, left panel) the market is apparently efficient, since the market price $p(k)$ —blue line in the figure—tracks the fundamental price $p_f(k)$. This is the result of the fact that our double-auction market, when determining the market




Fig 1. Box plots of the realizations of the herding intensity H (top panel) and the market efficiency \bar{e} (bottom panel) obtained with 100 simulations for each value of ζ .

<https://doi.org/10.1371/journal.pone.0239132.g001>




Fig 2. Dynamics of the asset price estimates performed by the traders for three sample simulations with the same fundamental price $p_f(k)$ (black line), and only differing in the sociability ζ . The dynamics of the market price $p(k)$ are depicted in blue.

<https://doi.org/10.1371/journal.pone.0239132.g002>

price, tends to average the price expectations of the different agents, which range around $p_f(t)$ in a band of width 20. Note that increasing ζ has a twofold effect on the market dynamics. Indeed, it both reduces the variance of the expected fundamental price \hat{p}_i across the agents, and steers the \hat{p}_i -s towards the expectation of the most influential (wealthy) nodes. For moderately high values of ζ (see the central panel of Fig 2), this improves market efficiency, moving the market price closer to the fundamental price, while for higher values of ζ (Fig 2, right panel), this can produce strong deviations of $p(k)$ from $p_f(k)$, thus generating market bubbles (where the market price differs more than 20% from the fundamental price) and dramatically reducing efficiency.

Controlling market efficiency

The analyses performed in the previous section illustrates how an increase of the herding intensity is not necessarily detrimental for market efficiency. However, when all the agents tend to conform their expectations to those of a group of influential leaders, efficiency is lost and market bubbles may emerge. In social trading, the market platform (as, for instance, *eToro* and *Zulutrade* [15]) provides the infrastructure to display profiles of the traders to other users. Therefore, the platform can manipulate the topology of interconnection and the kind of information shared among the investors. In what follows, we propose and numerically validate two alternative control strategies to enhance market efficiency, which can be translated into implementable market policies. Both strategies aim at avoiding that the opinion of a few determined the expectation of the group. The first control strategy, denoted *isolating the hubs*, prescribes that the strategies of the most influential nodes are not spread through the platform, while the second approach, denoted *anonymization*, consists in not disclosing the identity of the agents across the network. To test the effectiveness of the two policies, in the following we refer to the same numerical setup used to test the properties of the model in the absence of control.

Isolating the hubs. This control strategy leverages the vulnerability of scale-free networks to targeted attacks, which may disconnect the network by only removing the 2% of the most connected nodes [43]. Indeed, our market policy consists in selecting a subset composed of the most connected (the wealthiest) agents in the network, identified in the network as the nodes with the highest out-degree, and neutralize their influence by removing all their outgoing edges.




Fig 3. Box plots of the realizations of the herding intensity H (top panel) and the market efficiency \bar{e} (bottom panel) with 100 simulations for each value of δ . The green-shaded area corresponds to the box obtained for the case $\zeta = 0$.

<https://doi.org/10.1371/journal.pone.0239132.g003>

To assess the effectiveness of this approach, we focus on the case in which the uncontrolled market is less efficient, that is, when $\zeta = 1$. Then, we start to control (i.e. isolate) an increasing percentage δ (from 1 to 10, with step 1) of the most connected nodes. From Fig 3, we notice that, as the number of isolated nodes increases, the market regains efficiency. For each value of δ , 100 simulations were run, and a one-way ANOVA has been performed to test whether the herding intensity and market efficiency were significantly different from the case $\zeta = 0$, and we always obtained a p -value smaller than 0.005 (except the market inefficiency when the 3% of the nodes are isolated). We observe that, as δ increases, while the herding intensity remains substantially higher than when sociability is naught, market inefficiency becomes smaller than the case $\zeta = 0$ for $\delta > 5$ (see the green bands in Fig 3), further supporting the thesis that herding is not necessarily detrimental.

Anonymization. The second control action we propose is to anonymize the data exchanged among the traders. More specifically, each investor will receive from her neighbors only the expected return of the traded asset, but will not be aware of her identity. Therefore, she will weight equally the opinions of agents with heterogeneous wealth. When sociability is




Fig 4. Box plots of the realizations of the herding intensity H (top panel) and the market efficiency \bar{e} (bottom panel) obtained with 100 simulations of our model for each value of ζ when data are anonymized. The green-shaded area corresponds to the box obtained for the case $\zeta = 0$. The gray-shaded boxes replicate those reported in Fig 1 in the absence of anonymization.

<https://doi.org/10.1371/journal.pone.0239132.g004>

high, this mechanism will avoid the pathological condition in which all the agents follow the expectation of a single trader, but they will rather follow the *wisdom of the crowd* and the market price will tend to converge towards the average of the individual price estimation, which can be considered as a good approximation of the fundamental price. As illustrated in Fig 4, for all values of sociability, the anonymization strategy only mildly affects the herding intensity (see the comparison with the grey boxes in the top-panel). On the other hand, the anonymization policy makes sociality always beneficial for market efficiency. Indeed, for all $\zeta \neq 0$, the average inefficiency \bar{e} is always reduced if compared with the case of absence of sociality (see the green band in the bottom panel), and this difference is statistically significant for all $\zeta > 0.2$ (one-way ANOVAs, $p < 10^{-4}$).

By comparing the box plots in the absence and in the presence of the anonymization strategy (see the comparison with the grey boxes in the bottom panel of Fig 4), we observe that, for low values of the sociability, the choice of not disclosing the wealth has a negligible impact of market efficiency. Indeed, the reduced level of sociality avoids the prevalence of opinions of

one agent on the rest of the traders even when the wealth is disclosed. However, when $\zeta \geq 0.5$, anonymizing the data becomes crucial to preserve market efficiency, and a one-way ANOVA comparing the market with and without anonymization confirms a significant difference ($p < 10^{-4}$).

Robustness to variations in the risk aversion

Here, we evaluate the robustness of our findings to variations in the risk aversion of the traders. Specifically, we evaluated the effect of sociability for 10 selected values of the maximum risk attitude ϱ_{\max} , equally spaced on a logarithmic scale between 10^{-3} and 0.1 (ϱ_{\max} was set to 5×10^{-3} in Table 1). Values of ϱ_{\max} beyond 0.1 are not considered since they would correspond to excessively risk averse agents, thus yielding a market where no-one is willing to buy.

We notice that, in the absence of a policy to improve market efficiency, when the agents conform to the opinion of the most influential traders, market efficiency worsens, independent of the risk aversion of the agents, see Fig 5a. For all ϱ_{\max} smaller than 0.01, we identify a tipping point in the sociability ζ where we observe the transition from crowd wisdom to herding. The presence of a tipping point from beneficial to detrimental sociability is consistent with the findings in [44], where the authors highlight that this transition happens when the opinion of the agents conforms to that of a group of a few *experts*, which in our work would correspond to the wealthiest (i.e., most successful) traders. When ϱ_{\max} increases, the tipping point becomes smaller and smaller, and, when the agents are highly averse toward risk ($\varrho_{\max} > 0.01$), sociability can only increase inefficiency.

The effectiveness of the proposed policies to limit market inefficiency is confirmed both for the *isolating the hubs* and the *anonymization* strategy, independent of the risk aversion. Indeed, a comparison of Panels a and c in Fig 5 shows how the median market efficiency is always higher when the identity of the traders is not disclosed. Similarly, increasing the percentage δ of isolated hubs always yields an increase in market efficiency, see Fig 5b. Note also that the proposed policies also reduce the prevalence of the white color in all panels, which corresponds to a market where no trades occur, since the agents are too averse to risk.




Fig 5. In all panels, the colorbar corresponds to the median \bar{e} of the market inefficiency \bar{e} , obtained from 100 simulations for each parameter combinations. Panels a and c depicts \bar{e} as a function of the sociability ζ and of the maximal risk aversion ϱ_{\max} with or without anonymization, respectively, while Panel b refers to the strategy *isolating the hubs* and depicts \bar{e} as a function of the percentage δ of isolated nodes and of ϱ_{\max} . The dashed horizontal line corresponds to the value ϱ_{\max} selected in Table 1.

<https://doi.org/10.1371/journal.pone.0239132.g005>

Discussion

Leveraging a recently proposed artificial double-auction market, the analyses performed in this manuscript contribute a possible solution of the apparent dichotomy between herding and wisdom of the crowd. Indeed, although the presence of imitative behaviors in financial market is indisputable, contrasting empirical observations suggested that they can both be beneficial (hence called *crowd wisdom* or *rational adaptation*) or detrimental (hence called *herding*) for the market efficiency. Our numerical analyses illustrate that, when social interaction among the agents is present but not predominant, imitative behaviors can be beneficial. Indeed, each agent still considers her own estimation of fundamental price p_f but, combining it with the opinion of her neighbors, the variance of her price expectations shrinks, thus reducing the fluctuations around p_f . When social expectation dominates the individual estimations, the price expectations of all the traders will tend to coincide. However, since the agent will weight differently the neighbors' opinion depending on their wealth, their expectations will conform to those of few wealthier agents, similarly to what is observed in weighted and/or signed consensus protocols in opinion dynamics [45–47]. As a result, when the wealthier agents misestimate the fundamental price, the other agents blindly follow, and market bubbles may emerge. In a different context, a similar effect of sociality has been illustrated in the study by Lorenz *et al.* [48], where a group of participants had to solve different estimation tasks about geographical facts and crime statistics. The authors showed that the estimation accuracy can be undermined by sociality. As in our financial market model, a moderate advantage may only persist when sociality is modest, while the accuracy of the estimation substantially worsens when the crowd follows the opinion of a few highly confident individuals.

Considering the impact that sociality has on price estimation, we developed two alternative strategies to contain market inefficiency also when the social interaction among the agents is prevalent. In particular, we assumed that the market institution can manipulate the network topology in order to prevent instabilities. This is the case, for instance, of social trading platforms that may decide the kind of information shared among the agent, and to appropriately tailor the interaction topology. The first proposed control strategy aims at limiting the influence of a small fraction of the network nodes by eliminating their outgoing edges. Translated into a policy for the social trading platforms, this consists in not disclosing the trading strategies of the wealthiest agents. Our results showed that it suffices to adopt this policy on the wealthiest 3% of the traders to avoid the prevalence of the expectations of a few on the market sentiment, thus regaining efficiency. A second alternative control strategy consists instead in anonymizing the expectations exchanged among the agents. This means that the platform forwards to the out-neighbors of a node only her expectation, but not her wealth. In this way, every agent is going to equally account for the expectations of all her neighbors. The efficacy of this egalitarian mechanism is that it filters the most extreme expectations. In this way, the effect of sociality is solely beneficial, since it both reduces the variance of the price expectations, and centers their mean on the fundamental price. There are several directions along which our research can be extended. Since previous work suggested that herding may be more prominent in a multi-asset market [49], the efficacy of the proposed strategy in containing market inefficiency should also be tested when one or more assets are added in the market. Furthermore, experimental observations from world-wide financial markets and recent numerical studies have reported the possible emergence of multiple clusters of herding [50, 51]. Extending the model at a larger scale could be instrumental to identifying how the behavioral and topological features of the market may combine to determine clustered trading patterns.

Author Contributions

Conceptualization: Fabio Della Rossa, Lorenzo Giannini, Pietro DeLellis.

Data curation: Fabio Della Rossa.

Formal analysis: Fabio Della Rossa, Pietro DeLellis.

Funding acquisition: Pietro DeLellis.

Investigation: Fabio Della Rossa, Lorenzo Giannini, Pietro DeLellis.

Methodology: Fabio Della Rossa, Pietro DeLellis.

Project administration: Fabio Della Rossa, Pietro DeLellis.

Software: Fabio Della Rossa.

Supervision: Fabio Della Rossa, Pietro DeLellis.

Validation: Fabio Della Rossa, Pietro DeLellis.

Visualization: Fabio Della Rossa, Pietro DeLellis.

Writing – original draft: Fabio Della Rossa, Pietro DeLellis.

Writing – review & editing: Fabio Della Rossa, Pietro DeLellis.

References

1. Fama EF. Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance*. 1970; 25(2):383–417. <https://doi.org/10.1111/j.1540-6261.1970.tb00518.x>
2. Weintraub ER. Neoclassical economics. *The Concise Encyclopedia of Economics*. 2002;.
3. Colander D. The death of neoclassical economics. *Journal of the History of Economic Thought*. 2000; 22(2):127–143. <https://doi.org/10.1080/10427710050025330>
4. Tversky A, Kahneman D. Prospect theory: an analysis of decision under risk. *Econometrica*. 1979; 47(2):263–292. <https://doi.org/10.2307/1914185>
5. Shiller RJ. From efficient markets theory to behavioral finance. *Journal of Economic Perspectives*. 2003; 17(1):83–104. <https://doi.org/10.1257/089533003321164967>
6. Chiarella C, He XZ, Hommes C. Moving average rules as a source of market instability. *Physica A: Statistical Mechanics and its Applications*. 2006; 370(1):12–17. <https://doi.org/10.1016/j.physa.2006.04.026>
7. Teeter P, Sandberg J. Cracking the enigma of asset bubbles with narratives. *Strategic Organization*. 2017; 15(1):91–99. <https://doi.org/10.1177/1476127016629880>
8. Smith VL, Suchanek GL, Williams AW. Bubbles, crashes, and endogenous expectations in experimental spot asset markets. *Econometrica: Journal of the Econometric Society*. 1988;p. 1119–1151.
9. Lei V, Noussair CN, Plott CR. Nonspeculative bubbles in experimental Asset markets: Lack of common knowledge of rationality vs. actual irrationality. *Econometrica*. 2001; 69(4):831–859. <https://doi.org/10.1111/1468-0262.00222>
10. Berg N, Gigerenzer G. As-if behavioral economics: Neoclassical economics in disguise? *History of Economic Ideas*. 2010;p. 133–165.
11. Cocco L, Concas G, Marchesi M. Using an artificial financial market for studying a cryptocurrency market. *Journal of Economic Interaction and Coordination*. 2017; 12(2):34–365.
12. Tedeschi G, Iori G, Gallegati M. The role of communication and imitation in limit order markets. *The European Physical Journal B*. 2009; 71(4):489–497. <https://doi.org/10.1140/epjb/e2009-00337-6>
13. Cavalli F, Naimzada A, Pireddu M. An evolutive financial market model with animal spirits: imitation and endogenous beliefs. *Journal of Evolutionary Economics*. 2017; 27:1007–1040. <https://doi.org/10.1007/s00191-017-0506-8>
14. Biondo AE. Information versus imitation in a real-time agent-based model of financial markets. *Journal of Economic Interaction and Coordination*. 2019;p. 1–19.

15. Glaser F, Risius M. Effects of transparency: analyzing social biases on trader performance in social trading. *Journal of Information Technology*. 2018; 33(1):19–30. <https://doi.org/10.1057/s41265-016-0028-0>
16. Tedeschi G, Iori G, Gallegati M. Herding Effects in Order Driven Markets: The Rise and Fall of Gurus. *Journal of Economic Behavior & Organization*. 2012 1; 81.
17. Cont R, Bouchaud JP. Herd behavior and aggregate fluctuations in financial markets. *Macroeconomic Dynamics*. 2000; 4(2):170–196. <https://doi.org/10.1017/S1365100500015029>
18. Haghani M, Cristiani E, Bode NW, Boltes M, Corbetta A. Panic, irrationality, and herding: three ambiguous terms in crowd dynamics research. *Journal of Advanced Transportation*. 2019; 2019.
19. Haghani M, Sarvi M. Following the crowd or avoiding it? Empirical investigation of imitative behaviour in emergency escape of human crowds. *Animal Behaviour*. 2017; 124:47–56. <https://doi.org/10.1016/j.anbehav.2016.11.024>
20. Haghani M, Sarvi M. Imitative (herd) behaviour in direction decision-making hinders efficiency of crowd evacuation processes. *Safety Science*. 2019; 114:49–60. <https://doi.org/10.1016/j.ssci.2018.12.026>
21. De Long JB, Shleifer A, Summers LH, Waldmann RJ. Positive feedback investment strategies and destabilizing rational speculation. *the Journal of Finance*. 1990; 45(2):379–395. <https://doi.org/10.1111/j.1540-6261.1990.tb03695.x>
22. Blasco N, Corredor P, Ferreruela S. Does herding affect volatility? Implications for the Spanish stock market. *Quantitative Finance*. 2012; 12(2):311–327. <https://doi.org/10.1080/14697688.2010.516766>
23. Wolfers J, Zitzewitz E. Prediction markets. *Journal of Economic Perspectives*. 2004; 18(2):107–126. <https://doi.org/10.1257/0895330041371321>
24. Peeters T. Testing the Wisdom of Crowds in the field: Transfermarkt valuations and international soccer results. *International Journal of Forecasting*. 2018; 34(1):17–29. <https://doi.org/10.1016/j.ijforecast.2017.08.002>
25. Bottazzi G, Giachini D. Far from the madding crowd: collective wisdom in prediction markets. *Quantitative Finance*. 2019; 19(9):1461–1471. <https://doi.org/10.1080/14697688.2019.1622285>
26. DeLellis P, DiMeglio A, Garofalo F, Lo Iudice F. The evolving cobweb of relations among partially rational investors. *PloS One*. 2017; 12(2):e0171891. <https://doi.org/10.1371/journal.pone.0171891>
27. Giannini L, Della Rossa F, DeLellis P. A Partially Rational Model for Financial Markets: The Role of Social Interactions on Herding and Market Inefficiency. In: Cherifi H, Gaito S, Mendes JF, Moro E, Rocha LM, editors. *Complex Networks and Their Applications VIII*. Cham: Springer International Publishing; 2020. p. 535–546.
28. Friedman D. The double auction market institution: A survey. *The double auction market: Institutions, theories, and evidence*. 1993; 14:3–25.
29. Bouchaud JP. Price impact. *Encyclopedia of Quantitative Finance*. 2010;. <https://doi.org/10.1002/9780470061602.eqf18006>
30. Chiarella C, Iori G, Perelló J. The impact of heterogeneous trading rules on the limit order book and order flows. *Journal of Economic Dynamics and Control*. 2009; 33(3):525–537. <https://doi.org/10.1016/j.jedc.2008.08.001>
31. Lux T, Marchesi M. Volatility clustering in financial markets: a microsimulation of interacting agents. *International Journal of Theoretical and Applied Finance*. 2000; 3(04):675–702. <https://doi.org/10.1142/S0219024900000826>
32. Stiglitz JE. Behavior towards risk with many commodities. *Econometrica: Journal of the Econometric Society*. 1969;p. 660–667.
33. Chiarella C, He X, et al. Asset price and wealth dynamics under heterogeneous expectations. *Quantitative Finance*. 2001; 1(5):509–526. <https://doi.org/10.1088/1469-7688/1/5/303>
34. Bartolozzi M. A multi agent model for the limit order book dynamics. *The European Physical Journal B*. 2010; 78(2):265–273. <https://doi.org/10.1140/epjb/e2010-10406-4>
35. Bouchaud JP, Mézard M, Potters M, et al. Statistical properties of stock order books: empirical results and models. *Quantitative Finance*. 2002; 2(4):251–256. <https://doi.org/10.1088/1469-7688/2/4/301>
36. Barabási AL, Albert R. Emergence of scaling in random networks. *Science*. 1999; 286(5439):509–512. <https://doi.org/10.1126/science.286.5439.509>
37. León C, Berndsen RJ. Rethinking financial stability: challenges arising from financial networks' modular scale-free architecture. *Journal of Financial Stability*. 2014; 15:241–256. <https://doi.org/10.1016/j.jfs.2014.10.006>
38. Tseng JJ, Li SP, Chen SH, Wang SC. Emergence of scale-free networks in markets. *Advances in Complex Systems*. 2009; 12(01):87–97. <https://doi.org/10.1142/S021952590900209X>

39. Ganesh A, Massoulié L, Towsley D. The effect of network topology on the spread of epidemics. In: Proceedings IEEE 24th Annual Joint Conference of the IEEE Computer and Communications Societies.. vol. 2. IEEE; 2005. p. 1455–1466.
40. Khanin R, Wit E. How scale-free are biological networks. *Journal of computational biology*. 2006; 13(3):810–818. <https://doi.org/10.1089/cmb.2006.13.810>
41. Lilliefors HW. On the Kolmogorov-Smirnov test for normality with mean and variance unknown. *Journal of the American Statistical Association*. 1967; 62(318):399–402. <https://doi.org/10.1080/01621459.1967.10482916>
42. Bartlett MS. Properties of sufficiency and statistical tests. *Proceedings of the Royal Society of London Series A-Mathematical and Physical Sciences*. 1937; 160(901):268–282.
43. Crucitti P, Latora V, Marchiori M, Rapisarda A. Efficiency of scale-free networks: error and attack tolerance. *Physica A: Statistical Mechanics and its Applications*. 2003; 320:622–642. [https://doi.org/10.1016/S0378-4371\(02\)01545-5](https://doi.org/10.1016/S0378-4371(02)01545-5)
44. Moussaïd M, Kämmer JE, Analytis PP, Neth H. Social influence and the collective dynamics of opinion formation. *PloS One*. 2013; 8(11).
45. Ren W, Beard RW, Atkins EM. A survey of consensus problems in multi-agent coordination. In: Proceedings of the 2005, American Control Conference, 2005. IEEE; 2005. p. 1859–1864.
46. Altafini C. Consensus problems on networks with antagonistic interactions. *IEEE Transactions on Automatic Control*. 2012; 58(4):935–946.
47. Proskurnikov AV, Matveev AS, Cao M. Opinion dynamics in social networks with hostile camps: Consensus vs. polarization. *IEEE Transactions on Automatic Control*. 2015; 61(6):1524–1536.
48. Lorenz J, Rauhut H, Schweitzer F, Helbing D. How social influence can undermine the wisdom of crowd effect. *Proceedings of the National Academy of Sciences*. 2011; 108(22):9020–9025. <https://doi.org/10.1073/pnas.1008636108>
49. Manahov V, Hudson R. Herd behaviour experimental testing in laboratory artificial stock market settings. *Behavioural foundations of stylised facts of financial returns*. *Physica A: Statistical Mechanics and its Applications*. 2013; 392(19):4351–4372. <https://doi.org/10.1016/j.physa.2013.05.029>
50. Chen T. Country herding in the global market. *Journal of Behavioral Finance*. 2020; 21(2):174–185. <https://doi.org/10.1080/15427560.2019.1663852>
51. Caetano MAL, Yoneyama T. A model for the contagion and herding. *Physica A: Statistical Mechanics and its Applications*. 2011; 390(23-24):4445–4450. <https://doi.org/10.1016/j.physa.2011.07.039>