



Munich Personal RePEc Archive

Cowboying Stock Market Herds with Robot Traders

Galimberti, Jaqueson and Suhadolnik, Nicolas and Da Silva,
Sergio

2016

Online at <https://mpra.ub.uni-muenchen.de/71758/>
MPRA Paper No. 71758, posted 08 Jun 2016 08:20 UTC

Cowboying Stock Market Herds with Robot Traders

Jaqueson K. Galimberti*

ETH Zurich

Nicolas Suhadolnik

BRDE Regional Development Bank for Extreme South

Sergio Da Silva

Federal University of Santa Catarina

Abstract

One explanation for large stock market fluctuations is its tendency to herd behavior. We put forward an agent-based model where instabilities are the result of liquidity imbalances amplified by local interactions through imitation, and calibrate the model to match some key statistics of actual daily returns. We show that an “aggregate market-maker” type of liquidity injection is not successful in stabilizing prices due to the complex nature of the stock market. To offset liquidity shortages, we propose the use of locally triggered contrarian rules, and show that these mechanisms are effective in preventing extreme returns in our artificial stock market. JEL codes: C63, G02.

Keywords: herding, robot trading, financial regulation, agent-based model.

1 Introduction

Stock markets are complex dynamic systems where the interactions between their composing agents have a crucial role in determining aggregate outcomes. By fostering the emergence of collective conformity, such as fads and social manias, these interactions can propagate small deviations of individual behavior from the fundamental valuations to the overall market. The occurrence of large fluctuations in stock prices can therefore be attributed to, at least to some extent, the emergence of herd behavior. From this point of view, this paper offers two main contributions. First, we

*Corresponding author. E-mail: galimberti@kof.ethz.ch. Phone: +41 44 6328529. Fax: +41 44 6321218. Correspondence address: KOF Swiss Economic Institute, ETH Zurich, LEE G 116, Leonhardstrasse 21, 8092 Zurich, Switzerland.

propose an agent-based model that accounts for the role of imitation in individuals' decisions and is capable of generating the large swings observed in actual stock markets. Second, we evaluate the effectiveness of policy mechanisms aimed at preventing the occurrence of large fluctuations caused by herd behavior.

Stock market crashes can have harmful effects on economic activity. Large losses of wealth can induce lower levels of consumption, and abrupt changes in the cost of capital can lead to severe distortions in investment decisions. Macroeconomic policy authorities should therefore be interested in the prevention of such stock market collapses. But recognizing the complexity of stock markets poses daunting challenges for policy making. Attempts to stabilize stock markets with monetary policy and financial regulation raise several impracticalities. First, bubbles are hard to spot and, until very recently, central banks were not inclined to respond to developments in asset prices (see Blanchard et al., 2012, for an assessment of how this view might be changing). Second, the experience of the recent crisis suggests that financial markets tend to innovate around regulations and the nature of risk-taking changes as the financial system gets more sophisticated (Edey, 2009). Another example is given by the Chinese authorities role and response to the 2014-15 bubble and crash in the equity market, where “broad-ranging interventions [...] appear to have increased investor uncertainty about financial sector policies” (IMF, 2015). In this paper, we argue that these aspects are typical of large, dynamic and complex systems that can self-organize into a critical state where minor perturbations may give rise to instabilities of macroscopic scales (see, e.g., Scheinkman and Woodford, 1994; Bak and Paczuski, 1995). One key issue is that such extreme events cannot be predicted and therefore require a different paradigm for the design of effective interventions.

We set up an artificial stock market with autonomous agents that interact in a two-dimensional lattice using simple decision rules based on imitative and fundamentalist behavior. Imitation is a key component for the emergence of herding in our model, and it can be motivated from different theoretical reasons: Private information can generate incentives for a rational decision maker to follow others' actions in sequential environments (e.g., Banerjee, 1992; Bikhchandani et al., 1992, and Devenow and Welch, 1996; Bikhchandani and Sharma, 2001, for reviews of this literature). Imitation also can arise from social (Bernheim, 1994) and psychological factors that cause behavior to deviate from fully rational considerations (e.g., Kirman, 1993; Lux, 1995). Nevertheless, empirically, there is scarce evidence that distinguishes between these motives in real markets (see Cipriani and Guarino (2014) for a discussion on the disconnect between the empirical and theoretical literature). To circumvent this debate, imitative behavior is introduced in our model by the explicit assumption of a simple local interaction rule, rather in the spirit of Herbert Simon's bounded rationality (Simon, 1982).

As usual, prices in our model are determined at the market level in response to imbalances

between aggregate demand and supply for the stock. Hence, the ultimate cause of extreme returns in our model is oscillations of market liquidity, which is consistent with recent accounts of how the 2008-09 financial market turmoil propagated (see Brunnermeier, 2009, for a review). Nevertheless, the emergence of large liquidity imbalances is not implicit in our model behavioral assumptions, nor by its market micro-structure. It is the local interactions architecture of our model that amplify clustered shortages of liquidity and have a significant impact at the overall market level.

The inherent complexity that these interactions prompt often restrict the feasibility of standard analytical tools for realistic inferences. One solution is the use of the agent-based computational approach (see Tesfatsion and Judd, 2006), where dynamic systems of interacting agents are computationally modeled to facilitate generative explanations (Epstein, 2007). Because the estimation of agent-based models is complicated by the lack of simple analytical solutions (see Grazzini and Richiardi, 2015, and references therein), we developed an empirical strategy based on a goodness-of-fit measure for the whole distribution of stock returns generated by our model. Specifically, we calibrate our model to replicate actual stock markets data, which present distribution of returns characterized by heavy tails, i.e., extreme returns are more likely to occur than Gaussianity would imply. Using simulations, we then show that our model is capable of matching the empirical distribution of daily returns of the Dow Jones Industrial Average (DJIA) index from 1996 to 2012.

We then turn our focus to the design and evaluation of policy schemes aimed at the prevention of sudden liquidity dry-ups. To that end, we conduct several counter-factual exercises using our calibrated model. First, we show that an “aggregate market-maker” type of liquidity provision policy is only partially effective for the stabilization of our artificial stock market. More generally, we argue that a policy design that neglects the interconnections between individual decisions and their scaling up to aggregate outcomes can be misleading and costly.

To account for the complex nature of stock markets, we propose the use of a system of trading algorithms, or robot traders, as described in Suhadolnik et al. (2010). Robot traders have been around for years, but used only for private gain. More recently, their use for high-frequency trading has been the cause of intense debate on whether their effects are beneficial or harmful to the functioning of financial markets, and how regulation should cope with the rapid pace of their technological innovation (see the reviews by Foucault, 2012; Kirilenko and Lo, 2013, and Farmer and Skouras, 2013, for an ecological perspective). Here, instead, we propose their systematic use for the benefit of public policy making.

In contrast to the practice of responding to aggregate observations, our robot traders are triggered locally to follow a contrarian rule in order to prevent stampede reactions caused by herd behavior. To prevent financial imbalances we also introduce a self-regulatory mechanism to decrease the robot’s contrarian behavior in response to its individual financial position. Hence, every robot has the autonomy to trade on the basis of its local information, but is also bound by its own

track of transactions. Also, addressing one of the key criticisms to asset price targeting, the robots' intervention does not depend on assessments of the stock's fundamental value. The only requirement is the introduction of a parallel system of autonomous trading algorithms that will gather information in real-time at key junctures of the market structure.

There is an interesting parallel between our approach and the role of independent assessments in collective decision making. According to what is known as "Condorcet's jury theorem," named after its proponent, the 18th century French intellectual Marquis de Condorcet, the pooling of independent information held by multiple individuals can lead to better decisions than those relying on particular dictatorial assessments (see, e.g., Grofman et al., 1983; Young, 1988; Boland, 1989). In other terms, and nitpicking the popular belief that stock markets are sometimes driven by the madness of crowds, we devise a "crowd of robot traders" to restore the wisdom often associated with collective decisions (see Surowiecki, 2005; Landemore and Elster, 2012, for many examples).

The existence of interdependencies between the individual decisions, however, may lead to violations to the Condorcet's principle. When decisions are correlated, the effectiveness of information pooling through the majority rule tends to decrease (Ladha, 1992; Berg, 1993; Ladha, 1995). Clearly, this is the case in our crowd of robot traders – even though the robot traders are devised to operate autonomously, the interconnectedness of agents in our artificial market can give rise to correlated contrarian responses. To circumvent this issue, we further developed a coordination mechanism that splits the robots' action into two stages: First, the local information is collected and pooled by a financial policy authority. Next, the decisions of the robot traders are coordinated to take into account the general assessment of the market condition.

We find that the robot traders are capable of stabilizing the stock market and reshaping the distribution of returns towards a Gaussian distribution, while the self-regulatory mechanism guarantees its financial sustainability. Furthermore, with the aid of the coordination mechanism, the number of robots required to mitigate extreme events is substantially reduced, which means that our approach requires only tiny perturbations to the usual functioning of the stock market. The calibration of our model also evidenced some uncertainty regarding agent's sensitivity to the observation of a quorum in the local neighborhoods, and our results indicate the relevance of such a specification for the design of effective stabilization policies.

The remainder of this paper proceeds as follows: In Section 2, we present our artificial stock market model and describe the calibration approach we adopted to match statistical properties observed in the data. Section 3 describes the liquidity provision policies we considered in the counter-factual exercises of Section 4. Section 5 concludes the paper with some final remarks.

2 Complex Stock Market Model

The stock market is represented by a $L \times L$ square lattice where each cell, indexed by $i = 1, \dots, N (= L^2)$, represents an agent. Each agent holds a portfolio of money ($m_{i,t}$) and assets ($a_{i,t}$). For simplicity, we assume there is only one stock being traded and that there are no short sell constraints. Every period each agent has to decide whether she (or he) wants to buy or to sell one unit of the stock. Agents are assumed to make their decisions concurrently, after which a market clearing process takes place aggregating the individual demands and adjusting the price of the stock accordingly.

2.1 Behavioral Rules

We model the agents' decision-making process as a discrete choice, with the probabilities of each action determined symmetrically and by merging two investment strategies: an imitative rule and a fundamentalist rule. Because we assume symmetry, we can simplify the presentation by focusing on the determination of the probability that agent i makes a buy order at period t , $\pi_{i,t}^B$, which is given by

$$\pi_{i,t}^B = \omega_{i,t} \mathcal{I}_{i,t}^B + (1 - \omega_{i,t}) \mathcal{F}_{i,t}^B, \quad (1)$$

where $\mathcal{I}_{i,t}^B$ and $\mathcal{F}_{i,t}^B$ are the probabilities that the agent will choose to buy the stock based on imitation and fundamentals, respectively, and $\omega_{i,t}$ regulates the weight assigned to imitation.

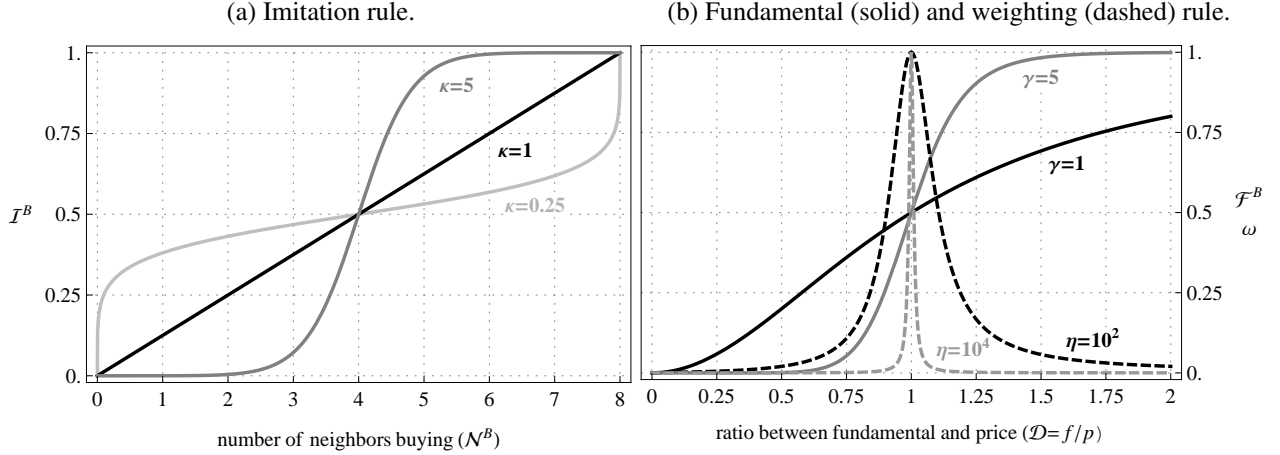
The imitative component is modeled according to a local interaction rule based on the agent's neighborhood, where the willingness to buy is an increasing function of the number of neighbors who have made a buy order in the previous period ($\mathcal{N}_{i,t-1}^B$), i.e.,

$$\mathcal{I}_{i,t}^B = \frac{(\mathcal{N}_{i,t-1}^B)^\kappa}{(\mathcal{N}_{i,t-1}^B)^\kappa + (\mathcal{N}_{i,t-1}^S)^\kappa}, \quad (2)$$

where $\kappa > 0$ controls the intensity of the response.

Inspired by the literature on consensual collective decision-making (see, e.g., Sumpter and Pratt, 2009), rule (2) displays a quorum-type response when $\kappa > 1$: the higher κ , the sharper is the increase in the probability of adopting a particular behavior once a quorum of agents performing that behavior is met in the neighborhood. Within the square lattice architecture, we assume agents interact with their eight surrounding neighbors, also known as Moore neighborhood; a quorum is given by a total of four neighbors adopting the same behavior. Furthermore, when $\kappa = 1$ the response becomes linear, and as κ decreases below 1 the response becomes insensitive to the quorum: the probability of imitation presents an initial sharp increase to a mid-range value once the first adopters are observed, but then remains around that level until a majority of adopters is

Figure 1: Behavioral Rules.



observed. These responses are illustrated in panel (a) of Figure 1.

In the context of the theoretical literature on imitation, mentioned in the introduction, our imitative rule is not derived from rationality assumptions. In particular, notice that (2) implies a naive assumption that the neighbors' previous actions fully reflect their private information, which neglects the correlation between the neighbors' own decisions on the basis of common sources (see Eyster and Rabin, 2010, 2014). Hence, the imitative component in our model follows in the spirit of models of bounded rationality (Simon, 1982).

For the fundamentalist rule, we assume that the agent holds a belief about the fair value of the stock, denoted by the fundamental value $f_{i,t-1}$, and weighs her willingness to buy or to sell the stock based on the gap between that evaluation and the stock's previous period price (p_{t-1}): when the price of the stock is below (above) its fundamental value, the agent will expect higher (lower) returns and therefore will increase her willingness to buy (sell) the stock. Letting $\mathcal{D}_{i,t-1}$ denote the ratio between $f_{i,t-1}$ and p_{t-1} , the probability that an agent will choose to buy based on fundamentals is given by

$$\mathcal{F}_{i,t}^B = \frac{\mathcal{D}_{i,t-1}^\gamma}{\mathcal{D}_{i,t-1}^\gamma + \mathcal{D}_{i,t-1}^{-\gamma}}, \quad (3)$$

where $\gamma > 0$ is a parameter regulating the agent's response to the deviations of the stock price from its perceived fundamental value. The different shapes of the responses obtained with this rule are illustrated in panel (b) of Figure 1.

The fundamental rule in (3) can be motivated from the literature on predictor selection under a discrete choice setup (Brock and Hommes, 1997), where γ represents the intensity of choice and measures how fast agents switch between different prediction strategies. Translated to our context, the interchange is between different investment strategies (buying low/selling high) that have corresponding prospects of return. Other than introducing standard asset pricing concerns

in our model, this fundamentalist component also has a key role in supporting the uncertainty required for herding behavior to emerge in a financial market (see Avery and Zemsky, 1998). We discuss this further below after introducing the price adjustment mechanism.

The fundamental beliefs are assumed to be determined exogenously to the model according to a log-normal distribution with a time-varying median¹, $\hat{\mu}_t$. In order to calibrate the model to match the trend swings observed in the index series of stock prices, we set $\hat{\mu}_t$ to correspond to a trend estimated from the data according to a procedure detailed in the next section. The dispersion of the heterogeneous beliefs across the agents is controlled by setting the variance of the log-normal distribution so that the variance of the fundamentals remains a constant fraction of the median fundamental value through time, i.e., $Var(f_{i,t}) = \sigma \hat{\mu}_t$. Using conventional notation, these distributional assumptions require that $\log[f_{i,t}] \sim N\left(\log[\hat{\mu}_t], \log\left[1/2 + \sqrt{1 + 4\sigma/\hat{\mu}_t}/2\right]\right)$. Notice that when $\sigma = 0$, agents hold identical fundamental beliefs, such that the only source of diversity in our model comes from their interactions within neighborhoods.

The deviation of the stock market price from its fundamental value also plays a role in the determination of the weights ascribed to each strategy. Here, we assume that the agents give more attention to the fundamental rule as the distance between the stock price and its fundamental value increases. Conversely, the agents will increase their reliance on the signal collected from their neighborhood as the fundamental price gap decreases. This behavior is captured by


$$\omega_{i,t} = \frac{1}{1 + \eta \tilde{\mathcal{D}}_{i,t-1}^2}, \quad (4)$$

where $\tilde{\mathcal{D}}_{i,t-1} = \log[\mathcal{D}_{i,t-1}]$, and $\eta > 0$ regulates the agents sensitivity to the relative deviations of the stock price from its fundamental value, as illustrated in panel (b) of Figure 1. From a statistical mechanics perspective, Rule (4) can be accountable for preventing the emergence of degenerate results in the dynamical system. In particular, the presence of positive feedbacks in the imitation strategy in the form of self-reinforcing collective behavior (see Sornette and Zhou, 2006) may favor the dominance of that rule in the determination of agents' final behavior.

Thus, our behavioral assumptions associate two main variables to the determination of agents' responses: the composition of their neighborhood and the deviation of the stock price from its fundamental value. The shape of these responses is illustrated in Figure 2 for both quorum-insensitive and quorum-sensitive imitative responses. Although π^B is strictly increasing in \mathcal{N}^B , the shape of the response depends on whether κ is smaller than or larger than 1. The response with respect to \mathcal{D} depends on two effects: the fundamental rule effect, making π^B strictly increase with \mathcal{D} ; and the

¹The median is the appropriate measure of central tendency in our model, where the number of agents buying and selling is the ultimate determinant of price adjustments, according to (5)-(6) below. Therefore, what matters for the "aggregate balance" of the fundamental rule is the number of agents with $f_{i,t}$ below and above $\hat{\mu}_t$, rather than the deviation of the mean of $f_{i,t}$ from $\hat{\mu}_t$.

Figure 2: Agents' Response Functions.



The surfaces represent the probability that an agent will attempt to buy (π^B) depending on the number of neighbors buying (N^B) and the ratio between the perceived fundamental value of the stock and its price (\mathcal{D}), according to (1)-(4).

weighting effect, which turns the relationship between π^B and \mathcal{D} negative when \mathcal{I}^B is sufficiently larger (smaller) than \mathcal{F}^B and $\mathcal{D} > 1$ ($\mathcal{D} < 1$).

2.2 Market Clearing

The model dynamics emerges as a result of the synchronous update of the agents' demands and a market clearing mechanism that randomly matches individual orders and adjusts the price of the stock according to an excess demand function. Letting

$$Z_t = \frac{N_t^B - N_t^S}{N}, \quad (5)$$

represent the (relative) excess demand for the stock. The price adjustment is modeled through a hyperbolic tangent functional form (see Plerou et al., 2002) according to

$$p_t = p_{t-1} (1 + \tanh [Z_t]). \quad (6)$$

This price adjustment process can be motivated as the action of a sluggish auctioneer who attempts to balance demand and supply for the stock (Chiarella et al., 2006; Lux, 2009). In spite of this underlying mechanism, we assume that market activity is generated solely by its composing traders. This is to say that, after the random matching between buyers and sellers is completed,

any remaining orders will be unsatisfied. Obviously, other than for regulating the market aggregate liquidity, this market clearing mechanism is also important for the record keeping of agents' transactions and the evolution of their portfolio, tracked by $m_{i,t}$ (money) and $a_{i,t}$ (assets).

When trade is sequential, Avery and Zemsky (1998) point out that the presence of a price adjustment mechanism may hamper the emergence of rational herding in financial markets. Because signals from the prices would temper the uncertainty that leads agents to act against their own private information, there must be multiple unknowns for herding to arise, e.g., agents must be uncertain about both the effect and the occurrence of information events on the value of the asset (see Park and Sabourian, 2011, for more general conditions). Although our model of herding is not explicitly based on rationality, and trade is simultaneous rather than sequential, the convolution between imitative and fundamentalist behavior in (1) also reflects one such case of multidimensional uncertainty. Particularly, our agents are uncertain about the fundamental value of the stock (Eq. 3) and about the quality of their neighbors' information (Eq. 4). Hence, our heuristic rule of imitation seems consistent with the requirements for the emergence of rational herding in sequential trading.

2.3 Data and Calibration

We generated artificial series of prices simulating the model with a square lattice of 100×100 agents. For every simulation, the lattice is initialized randomly with half the agents as buyers and the other half as sellers. As an empirical benchmark, we use the daily log returns of the Dow Jones Industrial Average (DJIA) index, corrected for inflation using the U.S. Consumer Price Index, from January 2, 1996 to December 31, 2012, consisting of a total of 4,277 observations². We also hold three years of data on the DJIA index, from 2013 to 2015, out of the calibration procedure in order to run an out-of-sample validation later on. To avoid sensitivity to initial settings, we ran our model for 4,400 periods, and therefore discarded the first 123 observations for the comparison between the simulated and empirical distributions.

Our model is purposely designed to simulate trading at daily or higher frequencies. In order to capture the lower frequencies of fluctuations, such as temporary trends commonly observed in series of stock price levels, we estimate the time-varying series of fundamentals from the data using the Hodrick-Prescott (HP) filter (Hodrick and Prescott, 1997). The HP filter is one of the most-used tools for the measurement of business cycles. It decomposes a series of observations into a trend and a cycle component as an approximation to a high-pass filter, where the maximum frequency of the cycles allowed to remain in the trend series is determined by a smoothing parameter, λ . For our

²The U.S. stock market is arguably among the most efficient and liquid stock markets nowadays, which leads to the question of how would the model work in less efficient markets, such as those of emerging economies. In fact, an earlier version of our model (Suhadolnik et al., 2010) provided a good adjustment to Brazilian stock market data. We leave a comparative analysis of our present model's adjustment to other markets for future research.

purposes we set $\lambda = 30,000$, a value that renders a trend that is stripped of cycles with frequencies up to approximately one quarter³.

The HP filter is also two-sided by design, which means that a point trend estimate is dependent on both lagged and leading observations of the original series. Hence, our approach is geared towards the view that the fundamental value reflects not only past prices information, but it is also a function of the future evaluations of the stock. One potential problem with this approach is that the symmetry of the filter is lost at sample endpoints (see, e.g., Galimberti and Moura, 2016), but here we deal with this issue by augmenting both ends of our sample with additional data; i.e., the fundamentals are obtained by applying the HP filter to data on the DJIA index from January 03, 1995 to December 31, 2015. The estimated trend, presented in panel (a) of Figure 3, is then introduced in our model to represent the deterministic portion, $\hat{\mu}_t$, of the exogenous series of fundamentals that agents hold for their daily evaluations.

To calibrate the model parameters, namely κ , γ , η , and σ , we ran a grid search procedure for several combinations of these parameters, attempting to minimize a measure of fitness to the data. For that purpose we adopt the two-samples Anderson-Darling (AD) goodness-of-fit statistic (see Scholz and Stephens, 1987), which compares the distribution of the simulated series of log returns to that obtained from the data. We also attempted to match some key statistical properties of the data, such as the returns auto-correlations, both in levels and absolute values to measure predictability and volatility clustering, respectively, and the kurtosis of the distribution of returns to measure the relevance of heavy tails. Table 1 presents statistics associated to some selected parametrizations⁴.

2.4 Model Dynamics

Overall, our results show that the best adjustment of the model to the data is found when the response of agents imitative rule is approximately linear. It was possible to find calibrations with a good fit for each of the three specifications of the imitative response, though very low/high values for κ were found to provide slightly poorer adjustments. It also was evident that as the quorum-sensitiveness of the agents' imitative response increased, the parameter on the intensity of choice in the fundamentalist rule had to be decreased to keep up with the adjustment of the model to the data. Thus, our model captures a trade-off between agents' sensitivity to their neighborhoods' information and their perceived deviations of the stock prices from its fundamental value, though the data and our calibration measures were not informative about this trade-off.


Our preferred specification is that with $\kappa = 0.75$, $\gamma = 2.50$, $\eta = 750$, and $\sigma = 0$, particularly

³The value of λ is calculated from the filter's frequency response function (see King and Rebelo, 1993) for the desired frequency at a 75 percent cut-off.


⁴Details of this calibration exercise are presented in Appendix A.1.

Figure 3: Time Evolution of Data and Calibrated Model Series.

(a) Price indexes and estimated fundamentals.



(b) Standardized returns.




The return series are standardized to have mean and standard deviation equal to zero and one, respectively. The model's series comes from the following parameter combination: $\kappa = 0.75$, $\gamma = 2.50$, $\eta = 750$, and $\sigma = 0$.

Table 1: Statistics on Series of Returns Series from Data and Simulated Model.

Series	Auto-correlations		AD test		Kurtosis	Tail exponent
	Returns	Abs.Rets.	Stat.	p-val.		
Dow Jones Industrial Average	-0.06	0.21	—	—	9.99	3.06
Model with quorum-insensitive imitative response:						
$\kappa = 0.25, \gamma = 2.50, \eta = 750, \sigma = 0$	-0.04	0.17	17.67	0.00	4.57	4.10
$\kappa = 0.75, \gamma = 2.50, \eta = 750, \sigma = 0$	-0.05	0.49	1.18	0.10	7.31	3.27
$\kappa = 0.75, \gamma = 2.50, \eta = 750, \sigma = 0.1$	-0.80	0.61	28.15	0.00	2.99	5.94
Model with linear imitative response:						
$\kappa = 1.00, \gamma = 1.00, \eta = 750, \sigma = 0$	0.47	0.38	43.70	0.00	2.79	6.24
$\kappa = 1.00, \gamma = 2.00, \eta = 750, \sigma = 0$	-0.06	0.54	0.66	0.18	5.98	3.39
$\kappa = 1.00, \gamma = 2.00, \eta = 750, \sigma = 0.1$	-0.64	0.37	27.55	0.00	3.00	5.91
Model with quorum-sensitive response:						
$\kappa = 3.00, \gamma = 1.00, \eta = 1,000, \sigma = 0$	0.14	0.67	48.69	0.00	5.51	3.23
$\kappa = 3.00, \gamma = 1.00, \eta = 15,000, \sigma = 0$	-0.06	0.79	0.69	0.17	6.93	3.35
$\kappa = 3.00, \gamma = 1.00, \eta = 15,000, \sigma = 0.1$	-0.38	0.13	28.07	0.00	2.99	5.88

Statistics are averages of 100 simulations of the model for each combination of parameters, only varying the random seed. The two-sample Anderson-Darling (AD) test compares the DJIA series of standardized log returns to those obtained from model simulations, and the median standardized statistic is reported. The null hypothesis is that both samples come from a common population and the p-values indicate the significance level at which this hypothesis can be rejected. The power law tail exponents are estimated according to Gabaix and Ibragimov (2011) focusing in the top 10% absolute returns.

Figure 4: Density Estimates of Data and Model Returns Distribution.



Densities were estimated using the Gaussian kernel smoothing function. The model's series comes from one typical simulation using the following parameter combination: $\kappa = 0.75$, $\gamma = 2.50$, $\eta = 750$, and $\sigma = 0$.


for presenting a kurtosis closer to that observed in the data. Also, using the AD test, we are not able to reject the null hypothesis that the simulated data and the actual data returns came from the same distribution at the 10 percent level of statistical significance. A visual assessment of the series of prices and log returns generated by this model specification is presented in Figure 3.

One remarkable characteristic of stock market returns is the presence of heavy tails, i.e., extreme events are more likely to occur than implied by a Gaussian distribution (Mandelbrot, 1963; Fama, 1965). This is evident in the data we used and our calibrated model seems to capture this property pretty well. Figure 4 shows that our model is able to capture the data deviations from Gaussianity at the tails of the return distribution up to ± 5 standard deviations. Moreover, our model is also in agreement with an established result of the econophysics literature regarding the decay of the probability of extreme returns; namely, the existence of an inverse cubic power law (see, e.g., Gopikrishnan et al., 1999) is confirmed by the estimates of tail exponents presented in Table 1.


When the fundamentals were allowed to vary across the agents, our calibration also showed a poor fit to the data in that it generated approximately Gaussian returns. Therefore, increasing agents' diversity of beliefs about the fundamental value of the stock is likely to dampen the destabilizing herding effects coming from their imitative behavior. Such a result is in contrast to a growing volume of literature where the diversity of beliefs is found to generate endogenous volatility in the absence of local interactions (see, e.g., Kurz, 1994; Brock and Hommes, 1998; Kurz et al., 2005; Branch and McGough, 2011). Hence, our analysis suggests that the effects of heterogeneous beliefs on volatility can be conditioned by the underlying assumptions on how agents interact, though we were not able to disentangle these features. We leave this issue open for

Figure 5: Boxplots of Model Statistics vs. Data Across Nonoverlapping Subsamples.

(a) Anderson-Darling goodness-of-fit statistic.



(b) Kurtosis of returns distribution.




The statistics come from 100 simulations of the model with $\kappa = 0.75$, $\gamma = 2.50$, $\eta = 750$, $\sigma = 0$. Outliers, i.e., returns larger/smaller than the whiskers, are not presented.

future research.

One potential concern with respect to the calibration of our model is that of data-snooping, i.e., the repeated use of the same reference data series for model selection purposes. In our case, the parameter sweep using the DJIA data series to calibrate the model may lead to an overfitting of the model to that series, particularly considering the large number of degrees of freedom in our imitation-based model. Nevertheless, it is important to recall that our approach is not dependent nor geared towards the model's predictive ability; instead, our focus is on matching some key statistical properties of the distribution of stock market returns, particularly those statistics capturing the incidence of instabilities.

To enhance our understanding of the fit of the model to the data we look at some adjustment statistics across nonoverlapping subperiods of the calibration sample (see Brock et al., 1992), as well as over an out-of-sample period from 2013 to 2015. Figure 5 presents the results focusing on the distributions of the AD statistic and the kurtosis of the returns distribution. Clearly, the statistics show that the fit of the model is not independent of the sample period; particularly, the model's ability to match the DJIA returns distribution tends to deteriorate in periods of reduced volatility, such as in the three years preceding the 2007-08 financial crisis. Whereas this observation is not surprising, given that our model is purposely designed to capture the emergence of periods of high volatility, we interpret these results as positive evidence that our approach is not severely affected by data-snooping.

Figure 6: Boxplots of Simulated Model Returns Before and After a Sharp Fundamental Shock.



The statistics come from 1,000 simulations of the model with $\kappa = 0.75$, $\gamma = 2.50$, $\eta = 750$, $\sigma = 0$, and $f_{i,t} = 1.50$ for $t = 1, \dots, 500$, and $f_{i,t} = 1.20$ afterwards. Outliers, i.e., returns larger/smaller than the whiskers, are not presented.

2.5 Role of Fundamentals

Due to our assumption of a smoothly time-varying series of fundamental values, the emergence of extreme returns in our model is mainly driven by the short-run effects of local interactions. That may give the mistaken impression that our model downplays the relevance of sharp fundamental reassessments in generating abrupt stock price changes, such as in mass sells triggered by unexpected poor performance of the firm, or a sudden deterioration of the macroeconomic conditions. In fact, an important feature of our model is its flexibility to account for these distinct sources of disturbances.

To clarify this point we simulate the model with a fictitious break in the fundamental value to see how the model behaves. Particularly, we run the model for 500 periods under the assumption of a constant fundamental value of 1.50, a value chosen to match the median of fundamentals used in the data calibration of the model above. We also adopt our preferred specification for the parameters $\kappa = 0.75$, $\gamma = 2.50$, $\eta = 750$, and $\sigma = 0$. At period 501 we then hit the model with an abrupt 20% negative shock to the fundamental, reducing it to 1.20. After repeating this simulation 1,000 times the resulting model dynamics is depicted in Figure 6, where we present the distributions of model returns over the periods that follow the fundamental shock.

We see that an abrupt change in the perceptions of the stock's fundamental value has an immediate impact on the model dynamics; in fact, the impact is initially stronger than the original shock, with a median return of about -60% after the abrupt reassessment of the fundamental; these effects

tend to fade away as time goes by and the market price resettles around the new fundamental value. Hence, the break in fundamentals is synergistically magnified by the herding effects, leading to an initial overshooting in the pricing of the stock. To conclude, an interesting feature of our model is the presence of a holistic interaction between fundamental assessments and local interactions that lead to the emergence of realistic stock market dynamics.

3 Liquidity Provision Policies

The emergence of extreme events in our stock market model is directly related to temporary shortages of liquidity for those agents willing to buy or sell their stock, i.e., imbalances in the aggregate market demand for the stock relative to its supply⁵. We now introduce and evaluate alternative policy schemes of liquidity provision aimed at offsetting these imbalances. Particularly, we look at the interesting and realistic case where the aggregate excess demand cannot be observed before the price adjustment process takes place. In these circumstances, the main challenge for policy design is to find a good prediction for the upcoming imbalance in the market.

3.1 Aggregate Market Maker

An obvious tentative solution to the problem of liquidity shortages is the introduction of an aggregate market maker. The idea is that by supplying the extra liquidity demanded by the market such a market maker may be able to prevent the occurrence of extreme and destabilizing returns in our stock market model.

With our timing assumption, the aggregate excess demand cannot be observed before the prices adjust. Hence, we suggest that such a market maker will respond to lagged measures of market activity. The aggregate excess demand in our model can be approximated as a linear function of the market clearing implied return⁶. Using such an approximation, our rationalization of the market-maker intervention as a response to perceived excess demand for the stock is given by

$$K_t = \{\Lambda r_{t-1} + K_{t-1}\}_{-\bar{K}}^{+\bar{K}}, \quad (7)$$

where K_t denotes the number of stocks supplied by the market maker at period t , r_{t-1} is the previous period log return of the stock, $\Lambda = (N + \bar{K})/2$, and $\{\bullet\}_{\min}^{\max}$ is a truncation operator used to impose a limit to the participation of the market maker. As a result, \bar{K} may be interpreted as a

⁵Within our model that is a direct implication of our assumption that stock prices are determined by a market clearing process that balances aggregate demand and supply for the stock (see Eq. 6).

⁶See Appendix A.2 for this derivation.

measure of the magnitude of the intervention. The presence of the lagged term K_{t-1} arises from the additional demand introduced by the market maker in the previous period.

3.2 Self-regulatory Robot Traders

The channel through which the liquidity imbalances arise in our model has, by design, a complex architecture due to the distributed effects of local interactions between the agents in this market. Thus, we argue that an aggregate intervention disregarding the complex nature of the market is of limited effectiveness for the stabilization of stock returns. To account for these local connections, we propose a novel self-regulatory scheme based on locally triggered automatic trading algorithms, or “robot traders” for short.

At any given period, the stock market model is now populated by two types of agents: human and robot traders, the latter indexed by $j = 1, \dots, Q$ and randomly distributed on top of the baseline model square lattice architecture of the market. We also assume that the robot traders are able to collect information about the last period decisions of the human agents within the neighborhood of their location. The distinctive feature of the policy rule of the robot traders is that they follow a contrarian rule relative to such information. Also in contrast to the human traders, the robots decisions are solved deterministically. Namely, robot j decides to make a buy order at period t if it observes that⁷

$$\left(\tilde{\mathcal{N}}_{j,t-1}^S\right)^{\beta[-a_{j,t}]} > \left(\tilde{\mathcal{N}}_{j,t-1}^B\right)^{\beta[a_{j,t}]}, \quad (8)$$

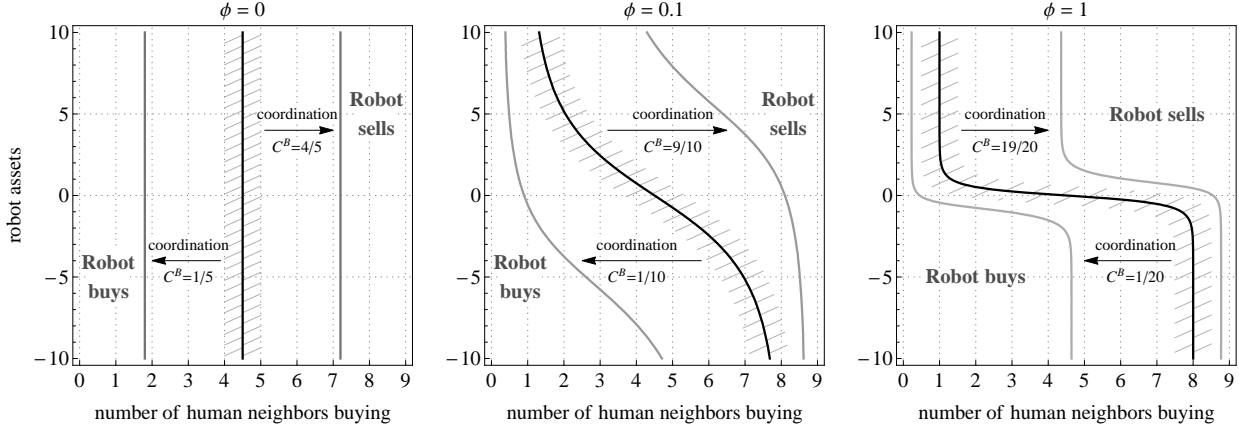
where $\tilde{\mathcal{N}}_{j,t-1}^S$ and $\tilde{\mathcal{N}}_{j,t-1}^B$ denote the number of humans in the robot’s neighborhood who have made a sell and a buy order in the previous period, respectively, and $\beta[\bullet]$ regulates the intensity of the contrarian response, which is assumed to be a function of the robot assets holdings. Particularly, we found that the functional form given by $\beta[x] = 1 + \tanh[\phi x]$ renders the effects of interest. When the robot’s holdings is off-balance, for example $a_{j,t}$ is large and positive (negative) due to successive buys (sells) of the stock in the previous periods, the contrarian response becomes less sensitive to the number of human sellers (buyers) in the neighborhood. The solid black schedules in Figure 7 illustrate how the robot’s decision is affected by its portfolio position for different values of ϕ . Logically, the robot places a sell order if condition (8) is not observed.

The mechanism underlying the robot traders’ decisions is intentionally designed to counteract the human agents tendency to herd. We advanced a similar rule in a previous work (Suhadolnik et al., 2010). Here, in contrast, we abandoned the unrealistic assumption that the human agents are directly influenced by the robots through imitative behavior⁸. This allows us to focus on the

⁷In Appendix A.3 we show that this contrarian rule is symmetric to the humans imitative rule.

⁸Results under the assumption of integrated robots are similar to what we find without this assumption. These are available upon request.

Figure 7: Contrarian Rule Responses.



The solid black schedules represent the combinations of \tilde{N}_j^B and a_j for which condition (8), or (9) with $\theta = 0$, of the contrarian rule is satisfied as an equality. The gray schedules represent cases with coordination, setting $\theta = 1/2$, for different levels of C^B , which stands for the fraction of robots deciding to buy in the interim stage.

liquidity provision effects that the robots have in the market prices, which in turn end up affecting agents' decisions through their fundamentalist concerns.

3.3 Coordinated Robots Intervention

Although the robot traders consist of a population of autonomous trading algorithms distributed across the market, there is still a unique financial policy authority behind their implementation⁹. That means there is some scope for communication and coordination among the robot traders. Namely, the financial policy authority can collect information about the market conditions that every robot can infer from their individual neighborhoods, and then alter each robot's action in a coordinated fashion. Such a coordinated intervention may also be viewed as a mixture of the previous approaches.

We modeled this intervention in a two-stage process. In the first stage, the ‘‘consultation round,’’ the robot traders solve condition (8) for an interim decision based on their corresponding neighborhood information. The financial authority then collects all these decisions so as to compute the fraction of robots that would be buying and selling without coordination, denoted by $C_t^B = Q_t^B/Q$ and $C_t^S = Q_t^S/Q$, respectively. In the second stage, this information is sent back to the robots in the form of a multiplier that will determine their final decision of buying (or selling otherwise)

⁹In the case of the United States, for example, such an authority could be represented by the Securities and Exchange Commission.

according to¹⁰

$$(1 - 2\theta C_t^S) \left(\tilde{\mathcal{N}}_{j,t-1}^S \right)^{\beta[-a_{j,t}]} > (1 - 2\theta C_t^B) \left(\tilde{\mathcal{N}}_{j,t-1}^B \right)^{\beta[a_{j,t}]}, \quad (9)$$

where $\theta \in (0, 1)$ is a parameter that controls the degree of interdependence between the robot traders. If $\theta = 0$, there is no coordinated action, and, as $\theta \rightarrow 1$, coordination among the robots grows in importance. Hence, in our model we attempt to find the right mix of independence and interdependence between the robot traders by tuning parameter θ .

Furthermore, there are two additional effects to notice when this coordination rule is activated. First, if $C_t^B < 1/2$, the multiplier on the left-hand side of (9) is smaller than the multiplier on the right-hand side, and vice versa for $C_t^B > 1/2$, which implies that coordination will push the individual robot's behavior towards the majority interim decision. This is illustrated by the gray schedules in Figure 7. Second, when $\theta = 1$, the right-hand side of (9) becomes negative if $C_t^B > 1/2$, whereas if $C_t^B < 1/2$, it is the left-hand side that becomes negative; i.e., in full coordination, all the robots will follow the majority interim decision, entirely disregarding their individual neighborhood information.

This idea of a coordinated intervention can be further motivated by the intriguing effects that communication has on the pooling of independent information in collective decision making. By creating interdependencies between decision makers, communication can facilitate information pooling (Ladha, 1995) at the same time that it can also cause the phenomenon known as “group-think” (Janis, 1982), where pressure to conform among the members of a group can narrow the range of opinions and lead to the emergence of informational cascades that amplify individual errors (Bikhchandani et al., 1992). A key issue in this context is to find the “right mix” of independence and coordination between the individuals in the group¹¹, which in our model is done by tuning θ .

3.4 Market Clearing Considerations

The policy interventions we have introduced in this section also require adjustments to the process of market clearing defined in the baseline model. To take into account the introduction of the aggregate market maker and the robot traders, the market aggregate excess demand is more generally defined as

¹⁰See Appendix A.4 for the origins of this rule as a linearly weighted response to the pooled and the individual pieces of information.

¹¹Nature appears to have solved this problem by endowing ants (Kirman, 1993) and bees (List et al., 2009) with simplistic rules.

$$\tilde{Z}_t = \frac{N_t^B - N_t^S - K_t + Q_t^B - Q_t^S}{N + \bar{K} + Q}. \quad (10)$$

The process of random matching between supply and demand in the market also needs to take into account the additional offers submitted by both the market maker and the robot traders. Therefore, we also need to track the portfolio of money and assets of each robot, say by $m_{j,t}$ and $a_{j,t}$, respectively, and of the market maker, which can be traced at an aggregate level by $m_{K,t}$ and $a_{K,t}$, respectively.

Furthermore, we assume anonymous trading in our artificial stock market, so that there are no asymmetries between the orders originated from the different agents in the market. This has two important implications. First, the chances that a buy (sell) order submitted during a bullish (bearish) period does not find a matching sell (buy) order in the market are the same irrespective of whether the order originated from a human agent, a robot or the market maker. Second, robots occasionally end up trading among themselves, which means that the intervention is artificially generating some extra market activity. In spite of this downside for the robot traders approach, anonymity provides an important safeguard against the possibility that the human agents uncover the participation of the contrarian algorithms and attempt to exploit it for their own gains.

4 Counter-factual Exercises

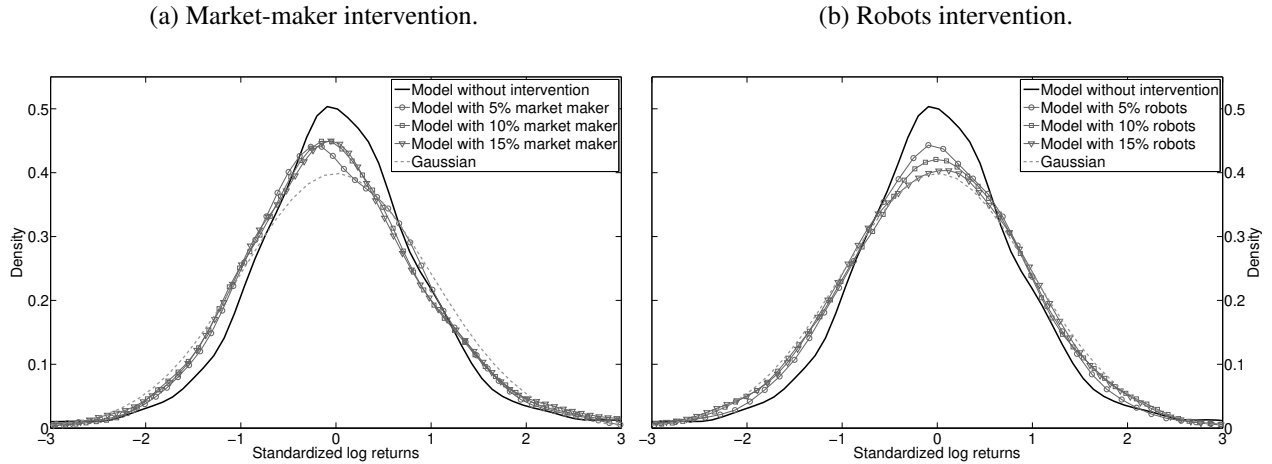
We now evaluate the effectiveness and the costs associated with the different interventions we have formulated in the previous section. In the spirit of a counter-factual exercise, we took our calibrated model as representative of how actual prices and returns are determined in a stock market and evaluated the effectiveness of the liquidity provision policies to prevent the occurrence of extreme events in this artificial market.

4.1 Visual Inspection

We begin with a visual inspection of the results from one typical simulation of our preferred specification of the model. In order to evaluate the stabilization effects of different policies, Figure 8 presents the distributions of returns associated with varying magnitudes of the interventions. Taking the Gaussian distribution as a reference, the goal was simply to flatten the distribution of the model-generated returns towards Gaussianity.

Clearly, the results in Figure 8 favor the robots intervention: panel (a) shows that although the aggregate market-maker intervention is partially successful in stabilizing returns, increasing the magnitude of the intervention brings no further improvements beyond that achieved with $\bar{K} = 500$; a different picture emerges for the robots' case, in panel (b), where a 15 percent intervention

Figure 8: Density Estimates of Model Returns Distribution With and Without Interventions.



Densities estimated using the Gaussian kernel smoothing function. The model's series come from the following combination of parameters: $\kappa = 0.75$, $\gamma = 2.50$, $\eta = 750$, and $\sigma = 0$. The magnitude of the interventions are denoted in relative terms to the number of human agents in the grid. The robots are self-regulated with $\phi = 0.1$.

is capable of bringing the model distribution of returns very close to the Gaussian distribution. Similar stabilization results are obtained when the robots are not self-regulated, i.e., $\phi = 0$.

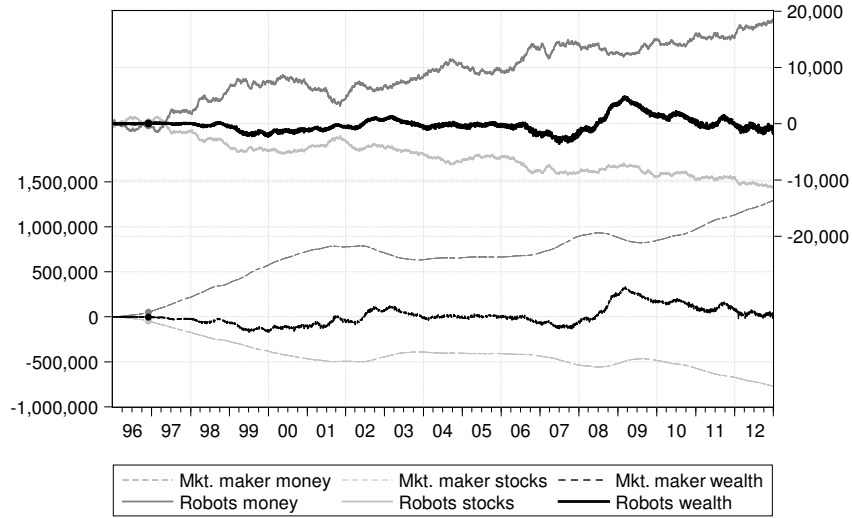
To assess the costs associated with the interventions, we look at the evolution of the aggregate financial position of these liquidity provision mechanisms. Particularly, we track the amounts of money and stocks accumulated by the purchases and sells of the interventions. Because we do not impose short selling constraints in our model, the participants of the artificial stock market can accumulate negative positions either in money or in stocks holdings. Nevertheless, from a policy making, practical standpoint, the sustained provision of liquidity to the market by selling stocks will require an equal amount of purchases in order to cover such positions. Hence, both the accumulation of a positive or negative balance of money by the intervention policy may be interpreted as a measure of its cost. To circumvent this issue, we also consider the wealth positions associated with each intervention, which are calculated taking into account the stock prices.

Figure 9 presents the evolution of the financial position of the interventions for the particular simulation we are considering in this section. There are three policy specifications to consider: the market maker, the robots without self-regulation ($\phi = 0$), and the robots with self-regulation ($\phi = 0.1$). The former two are plotted jointly in panel (a), and the latter is plotted in panel (b).

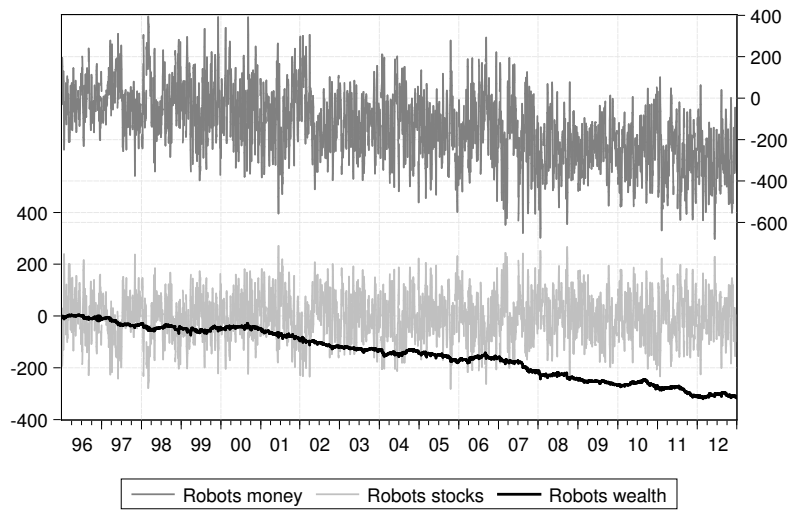
We reach two main conclusions from these results. First, the costs associated with the market-maker intervention are substantially higher than those incurred with the robots. Comparing the instances where both policies are free to operate without any self-regulation feature, as illustrated in panel (a), we calculated that the maximum (minimum) wealth of the market-maker intervention is about 70 (50) times larger than that of the robots.

Figure 9: Evolution of Intervention Costs.

(a) Market maker and robots without self-regulation.



(b) Self-regulatory robots.



The financial positions are obtained by simulating the model individually for each intervention, with the following parameter combination: $\kappa = 0.75$, $\gamma = 2.50$, $\eta = 750$, and $\sigma = 0$. The magnitude of the interventions are given by $Q = \bar{K} = 1,000$, and the self-regulated robots in panel (b) have $\phi = 0.1$.

Our second conclusion relates to the evident success achieved by our self-regulatory scheme to control the financial sustainability of the robots' operation, as illustrated in panel (b) of Figure 9. The self-regulatory robots incurred costs much smaller than their corresponding unregulated versions and the aggregate market maker. Clearly, this improvement resulted from the robots' ability to avoid the lingering financial imbalances at the individual level, while keeping up with the intervention effectiveness at the aggregate level.

4.2 Averaged Statistics and Robustness

Given the stochastic nature of our model, drawing conclusions on the basis of one particular simulation may be subject to the effects of random noise. We now extend our analysis to account for these effects by looking at averaged statistics related to the costs and effectiveness of the liquidity provision policies. For robustness purposes, we also consider alternative calibrations of the model with respect to the assumption on the agent's response to their neighborhood information. Namely, we consider the parameter combinations highlighted with a gray shade in Table 1, covering the cases of quorum-insensitive, linear, and quorum-sensitive responses.

Table 2 presents statistics related to the effectiveness of the interventions. Particularly, the kurtosis of the distribution of returns generated by the model is a measure that captures the essence of our stability analysis: the closer it gets to 3, the closer the distribution is to the desired Gaussian benchmark. In the quorum-insensitive responses, we can see that our previous conclusions are statistically confirmed, with the robots presenting important improvements as the magnitude of the intervention increases. But, this is not generally true under the alternative specifications of the agents' imitative response: There are some cases (gray shaded) where the intervention was not effective to reduce the return kurtosis. However, notice that even under these specifications, increasing the intensity of the intervention to 10 percent tends to bring the desired stabilization goals.

The statistical results for the market-maker intervention are again informative of its partial success. In terms of kurtosis reduction, we can see that a market maker restricted to operate at a maximum share of 1 percent of the market is better, on average, than the same quantity of robots. Besides, notice that the measure of volatility clustering ($AC\ Abs.$) tends to be lower under the market-maker intervention. However, the downside of this mechanism comes with the increase of its intensity, where we can see that there is no general improvement to its stabilization effectiveness. For the case of a market populated by quorum-sensitive agents, there is evidence that the market maker may even destabilize the market.

The statistical evaluation of the costs associated with the interventions is rather challenging. Following our previous discussion, we summarize the financial position of each policy by tracking

Table 2: Interventions Effectiveness.

Policies	Returns statistics, by agents imitative response								
	Quorum-insensitive			Linear			Quorum-sensitive		
	Kurt.	AC	AC Abs.	Kurt.	AC	AC Abs.	Kurt.	AC	AC Abs.
Without intervention	7.31	-0.05	0.49	5.98	-0.06	0.54	6.93	-0.06	0.79
Robots without self-regulation ($\phi = 0$):									
$Q = 100$	7.09	-0.00	0.48	6.66	0.00	0.56	6.78	-0.06	0.79
$Q = 500$	5.37	0.12	0.40	5.19	0.20	0.53	7.14	-0.07	0.75
$Q = 1000$	4.36	0.11	0.28	3.68	0.26	0.39	4.71	-0.05	0.63
Robots with self-regulation ($\phi = 0.1$):									
$Q = 100$	7.02	-0.00	0.48	6.42	-0.00	0.56	6.89	-0.06	0.79
$Q = 500$	5.37	0.11	0.41	5.32	0.20	0.54	7.35	-0.06	0.76
$Q = 1000$	4.26	0.12	0.29	3.56	0.26	0.40	4.99	-0.05	0.65
Market maker:									
$\bar{K} = 100$	4.84	-0.00	0.20	4.63	0.14	0.29	4.47	-0.16	0.44
$\bar{K} = 500$	5.09	-0.10	0.19	4.71	0.12	0.28	6.94	-0.39	0.44
$\bar{K} = 1000$	4.67	-0.29	0.21	4.33	-0.03	0.22	8.24	-0.47	0.42

Statistics are averages of 100 simulations of the model for each specification, according to Table 1, only varying the random seed. Kurt. stands for kurtosis, AC for autocorrelation, and AC Abs. for the autocorrelation of absolute returns.

Table 3: Interventions Costs.

Policies	Interventions wealth relative to volume, by agents' imitative response					
	Quorum-insensitive		Linear		Quorum-sensitive	
	Min. WV	Max. WV	Min. WV	Max. WV	Min. WV	Max. WV
Robots without self-regulation ($\phi = 0$):						
$Q = 100$	-0.04	0.09	-0.05	0.14	-0.06	0.33
$Q = 500$	-0.14	0.24	-0.17	0.38	-0.14	0.74
$Q = 1000$	-0.25	0.32	-0.29	0.40	-0.17	0.52
Robots with self-regulation ($\phi = 0.1$):						
$Q = 100$	-0.00	0.01	-0.00	0.02	-0.00	0.07
$Q = 500$	-0.01	0.00	-0.01	0.01	-0.00	0.23
$Q = 1000$	-0.05	0.00	-0.07	0.00	-0.00	0.19
Market maker:						
$\bar{K} = 100$	-0.30	0.73	-0.20	0.75	-0.21	0.81
$\bar{K} = 500$	-9.15	28.10	-8.02	24.22	-8.53	18.35
$\bar{K} = 1000$	-17.16	75.15	-15.98	68.81	-16.73	60.99


Statistics are averages of 100 simulations of the model for each specification, according to Table 1, only varying the random seed. WV stands for the ratio between wealth and volume of the corresponding intervention, and the statistics reported are averages of the minimum and maximum values of this ratio throughout each simulation.

the wealth associated with its money and stocks holdings. We then turn these measures into relative terms calculating ratios relative to the volume traded each period. Finally, we average the minimum and maximum values of this ratio over the simulations conducted for each model specification and intervention policy.

Statistics on these wealth/volume ratios (WV) are presented in Table 3. Clearly, we again confirmed our conclusions from the visual assessment of a typical simulation. First, the results show that the self-regulation of the robots intervention is successful in reducing the magnitudes of financial imbalances accumulated by the intervention in the three characterizations of the agents imitative responses. One interesting observation in this respect is that the self-regulatory robots tended to accrue profits under the assumption of quorum-sensitive responses. But apart from this case, the distortionary effects due to the self-regulatory robots over the agents' distribution of wealth remained below 10 percent of the traded volume in the market.

Finally, the results in Table 3 for the market maker show that there could be a very high cost associated with this intervention. Whereas in the case with a small intervention the distortions to the distribution of wealth between the agents can vary between -30 percent to about 80 percent of the traded volume, as the intervention intensifies, these distortions become astronomical.

Figure 10: Coordinated Intervention of 2 Percent Robots.



Statistics are averages of 100 simulations of the model for each specification, according to Table 1, only varying the random seed. The specification of the robots' intervention is given by $Q = 200$ and $\phi = 0.1$.

4.3 Results With Coordinated Interventions

Overall, our experiments indicate that introducing coordination reduces the number of robots required to achieve our previous stabilization results, though there is a greater sensitivity to the assumption of how agents react to their neighborhood. Figure 10 summarizes some of the results we obtain under coordination, focusing on the case with only 200 robots, or 2 percent of the number of humans in the baseline model. Clearly, there is some heterogeneity with respect to how the coordination affects the robots effectiveness.

For the cases of quorum-insensitive and linear-response agents, increasing the coordination parameter up to $\theta = 0.8$ brings gradual improvements to the stabilization of the stock returns. Beyond this threshold the coordination strategy may lead to undesirable effects, such as the generation of bimodal distributions. A different picture emerges for the case with quorum-sensitive agents: The benefits from coordination only start to appear for $\theta > 0.6$, meaning that a tendency towards the majority rule seems to work better under these particular circumstances. Hence, one key issue for the successful design of the coordinated approach is the identification of which behavioral assumption better characterizes how agents react to the information collected among their peers in actual stock markets.

4.4 Risk of Counter-Robots Developments

One potential problem with the approach of mixing the aggregate and the local interventions is that it may increase the risk that the human agents uncover the ongoing policy and use such information to systematically exploit the contrarian measures for private gain. Besides, unlike our public-service robots, software-based traders have been around for years but used for private gain. They have sometimes been considered as culprits for extreme moves in markets. Some observers blame the first generation of robots for the crash of 1987 as they were the tools for delivering so-called “portfolio insurance.” Currently, programs are far more sophisticated and responsible for billions of dollars traded every day, particularly in the form of high-frequency trades. In practice, because no one knows whether an order is placed by a robot rather than a human trader, the regulator can only react in the aftermath of extreme events.

Conventional regulatory frameworks such as this have to adapt to circumstances that are changing too fast for regulation to succeed, and the robots have something to do with this. Financial regulation implemented by the authorities are too conspicuous to succeed. It is in the nature of markets that they will tend to innovate around regulations, and the nature of risk taking will inevitably keep changing as financial systems get more sophisticated (Edey, 2009). An advantage of contrarian algorithms is their crypticness; by stealthily taking contrarian positions at key junctures in the movement of stock markets, the advantage of the stabilizing robots over conventional financial regulation is unmistakable.

Nevertheless, one cannot at first rule out counter-actions from the profit-seeking human traders, which may even lead to catastrophic outcomes such as a biological arms race for the development of the most sophisticated automated trading mechanism. Thus, it is important to devise further mechanisms that introduce uncertainty about the form and the timing of the robots interventions. One possibility would be to add extraneous variability to such policy measures, which could be accomplished by turning the determination of the robot’s trades stochastic. Another interesting extension would be to consider an evolutionary approach (see Evstigneev et al., 2009, for a review), where investment strategies are allowed to adapt, and selection and reproduction forces could be assigned to the maintenance of an appropriate degree of diversity. We leave these considerations for future research.

5 Concluding Remarks

We developed an agent-based model of a stock market that, in spite of endowing agents with simple behavioral rules, incorporates complex structures of local interactions that lead to the emergence of herd behavior. After calibrating the model, we have shown that it is capable of matching some

statistical properties observed in actual stock market data, such as the low degree of serial predictability; moderate levels of volatility clustering; and leptokurtic distributions of returns. One key feature of our model was that periods of market instabilities were generated as the result of liquidity imbalances amplified by the local interactions.

Our main contribution came with an attempt to stabilize the artificial stock market through the design of liquidity provision policies. Remarkably, we have shown that an intervention in response to the aggregate state of the market is only partially effective for stabilization and incurs high financial costs. We argued that the complex nature of the stock market needs to be taken into account, and accordingly we proposed the use of contrarian robot traders. These robots are spread through the market to collect local information on the market conditions and trade autonomously using a contrarian rule.

We showed that the robot traders can successfully offset periods of liquidity shortages and, as a result, are effective in keeping market volatility under control. We also devised a self-regulatory mechanism that prevents the costs associated with their implementation to become excessively burdensome. Additionally, we analyzed the case for a coordination mechanism where the robot intervention occurs in two stages: First, the robots collect the information from their neighborhoods and communicate it to a financial authority. Then, the authority aggregates this information and regulates a coordinated action. We found this additional mechanism to allow a substantial reduction in the number of robots required for an effective stabilization.

Therefore, we conclude that in spite of the problem of predicting the evolution of complex systems, extreme events caused by herd behavior may be potentially avoidable with the use of a locally triggered intervention strategy, together with an understanding of the underlying mechanisms of decision-making in stock markets. Of course, there are many open practical issues in the way between turning our proposal into a realistic intervention mechanism, e.g.: Who will conduct the robots implementation? Who will finance their operation costs? How to control their risks? At this stage we can only foresee that the actual implementation of our proposed mechanisms shall be subject to a great deal of experimentation to guide their actual design and feasibility.

A Appendix

A.1 Calibration Exercise

The final calibration was split into two stages.¹² In the first stage, we simulated the model with a gross grid of parameter combinations, and compared the associated evaluation statistics in order to learn the relevance of each parameter. Apart from the return auto-correlations and kurtosis mea-

¹²During the model design stage, massive simulations were conducted.

asures, we adopted an Anderson-Darling goodness-of-fit statistic as an evaluation statistic, which was adjusted for the possibility of ties (see Scholz and Stephens, 1987, Eq. 7). This statistic was used in two ways: in pairs, comparing the distribution of returns from each replication of the simulation against the data; and, in a multiple sample version where we tested whether the different replications of the simulation came from the same distribution in order to check for the role of randomness in the model. To capture the level of disorder in the overall system, we also adopted an averaged measure of the entropy in our two-dimensional lattice of agents, which we attempted to maximize in order to prevent the emergence of persistent clustering patterns. Following Wolfram (1983), we computed a block entropy, with the block corresponding to the neighborhood definition.

The greater uncertainty about the interval of parameter values relied on the determination of γ and η , which are both related to the model's sensitivity to the deviation of the stock price from its perceived fundamental value. Hence, we considered a broader range of values of these parameters. The results from this exercise are summarized in the box-plot diagrams in Figure 11, from which we can draw the following inferences:


Imitative response (κ): This parameter appeared to have small effects over the model's sensitivity to the other parameters. Most variation in the dispersion of each statistic occurred for values of $\kappa < 5$.

Fundamentalist response (γ): This parameter presented similar patterns within each statistic for $\gamma > 5$, with the fit of the model deteriorating for higher values of γ ; we also found that higher values of γ tended to cause degenerate results (the entire grid turning to the buy or the sell state) due to its exponential effect; particularly, notice that the box-plots on the AD statistic for the test that all replications came from the same distribution are below the rejection line only for $\gamma = 1$ and $\gamma = 2$. This indicates that higher values increase the effects of randomness in the model. For the auto-correlation of returns and for the averaged entropy, $\gamma = 1$ and $\gamma = 2$ are also closer to the desired targets of low predictability and high disorder in the lattice of agents.

Weighting rule sensitivity (η): Similar results were observed regarding the insensitivity of the distribution of each statistic for $\eta > 1,000$; focusing on the AD statistic and the return kurtosis we can readily rule out $\eta = 10$, because it turns the model results on these statistics insensitive to the other parameters. However, smaller η s tend to bring better results in terms of auto-correlation of returns and averaged entropy.

Diversity of fundamental beliefs (σ): Overall, the adjustment of the model tended to deteriorate for higher values of σ . In particular, the returns generated by the model tended to the Gaussian distribution as diversity increased. Our main conclusion is that the simplifying assump-

Figure 11: Box-plots of Key Statistics from Model Simulations.



The statistics come from 100 replications (varying the random seed) of each of the 900 combinations of parameter values, i.e., a total 90,000 simulations of the model. Our calibration aims to minimize the AD goodness-of-fit statistics (dotted lines represent critical values at the 25 percent significance level), to match the data in relation to return autocorrelation and kurtosis (dotted lines represent statistics from data), and to maximize averaged entropy.

tion, which is that there is no diversity of beliefs about the fundamentals, i.e., $\sigma = 0$, is critical for the emergence of heavy tails in the distribution of our model returns.

In the second stage of calibration, we constructed a finer grid of parameter combinations focused on the regions around the values that performed well in the first stage. Namely, we focused on the case of homogeneous fundamental beliefs, $\sigma = 0$, and evaluated the model for every combination of the following parameter values: $\kappa = \{0.10, 0.25, 0.50, 0.75, 1.00, 1.25, 1.50, 1.75, 2.00, 2.50, 3.00\}$; $\gamma = \{1.00, 1.25, 1.50, 1.75, 2.00, 2.25, 2.50, 2.75, 3.00\}$; and $\eta = \{500, 750, 1000, 1250, 1500, 1750, 2000, 2500, 3000\}$. This yielded a total of 1,386 parameter combinations, which, after multiplying by the 100 replications, resulted in a total of 138,600 simulations. The calibrations presented in Table 1 were selected from these results, which are also available upon request.

A.2 Market-maker Response

To invert the price adjustment function with respect to the excess demand, first notice that from (6) the stock market log returns are given by

$$r_t = \log [1 + \tanh [Z_t]].$$

Because $\tanh x = \frac{e^x - e^{-x}}{e^x + e^{-x}}$, this can be simplified to

$$r_t = \log 2 + 2Z_t - \log [1 + e^{2Z_t}],$$

or

$$e^{r_t} = \frac{2e^{2Z_t}}{1 + e^{2Z_t}},$$

which, after manipulation, leads to

$$e^{2Z_t} = \frac{e^{r_t}}{2 - e^{r_t}},$$

$$Z_t = \frac{r_t - \log [2 - e^{r_t}]}{2}. \quad (11)$$

Assuming r_t is small, so that $e^{r_t} \simeq 1 \Rightarrow \log [2 - e^{r_t}] \simeq 0$, we can approximate the excess demand implied by a given return by a linear function as advocated in the main text. Substituting for the modified excess demand of (10), and assuming there are no robots, i.e., $Q = Q_t^B = Q_t^S = 0$, the linearly approximated excess demand is given by

$$\frac{N_t^B - N_t^S - K_t}{N + \bar{K}} = \frac{1}{2}r_t,$$

which implies

$$N_t^B - N_t^S = \frac{(N + \bar{K})}{2} r_t + K_t. \quad (12)$$

The market maker is assumed to respond to lagged excess demand according to

$$K_t = \{N_{t-1}^B - N_{t-1}^S\}_{-\bar{K}}^{+\bar{K}}. \quad (13)$$

Substituting (12), lagged by one period, into (13) produces the final specification of the market maker response function given by (7).

A.3 Symmetry Between Robot and Human Rules

Assuming that the robots behave according to a counter-imitative rule symmetric to that of the humans, (2), the probabilities that robot j makes a buy/sell order at period t are given by

$$\mathcal{I}_{j,t}^B = \frac{\left(\tilde{\mathcal{N}}_{j,t-1}^S\right)^{\tilde{\kappa}}}{\left(\tilde{\mathcal{N}}_{j,t-1}^B\right)^{\tilde{\kappa}} + \left(\tilde{\mathcal{N}}_{j,t-1}^S\right)^{\tilde{\kappa}}}, \quad \mathcal{I}_{j,t}^S = \frac{\left(\tilde{\mathcal{N}}_{j,t-1}^B\right)^{\tilde{\kappa}}}{\left(\tilde{\mathcal{N}}_{j,t-1}^B\right)^{\tilde{\kappa}} + \left(\tilde{\mathcal{N}}_{j,t-1}^S\right)^{\tilde{\kappa}}}, \quad (14)$$

respectively. Because the robots make deterministic decisions, they would choose to buy if $\mathcal{I}_{j,t}^B > \mathcal{I}_{j,t}^S$, but because $\mathcal{I}_{j,t}^B \equiv 1 - \mathcal{I}_{j,t}^S$, robot j 's decision to buy simplifies to

$$\frac{\left(\tilde{\mathcal{N}}_{j,t-1}^S\right)^{\tilde{\kappa}}}{\left(\tilde{\mathcal{N}}_{j,t-1}^B\right)^{\tilde{\kappa}} + \left(\tilde{\mathcal{N}}_{j,t-1}^S\right)^{\tilde{\kappa}}} > \frac{1}{2},$$

$$\left(\tilde{\mathcal{N}}_{j,t-1}^S\right)^{\tilde{\kappa}} > \left(\tilde{\mathcal{N}}_{j,t-1}^B\right)^{\tilde{\kappa}},$$

which is identical to condition (8) after substituting the exponents for the self-regulatory mechanism.

A.4 Robots Coordination Through Pooling of Information

With the two-stages coordination approach, the robots have to decide on the basis of two pieces of information: 1) what they collect from their neighborhood, and 2) the signal they receive from the financial authority about the other robots' interim decisions. Assuming that the robots combine this information linearly, the probability that robot j will place a buy order at period t is given by

$$\pi_{j,t}^B = (1 - \theta) \mathcal{I}_{j,t}^B + \theta C_t^B, \quad \pi_{j,t}^S = (1 - \theta) \mathcal{I}_{j,t}^S + \theta C_t^S,$$

where $\mathcal{I}_{j,t}^B$ and $\mathcal{I}_{j,t}^S$ follows from (14), and the other terms are explained in the main text. Because the robots make a deterministic decision, they choose to buy if $\pi_{j,t}^B > \pi_{j,t}^S$, and sell otherwise. I.e., robot j 's condition to buy is given by

$$(1 - \theta) \mathcal{I}_{j,t}^B + \theta C_t^B > (1 - \theta) \mathcal{I}_{j,t}^S + \theta C_t^S, \quad (15)$$

but as $\mathcal{I}_{j,t}^B \equiv 1 - \mathcal{I}_{j,t}^S$,

$$(1 - \theta) (1 - 2\mathcal{I}_{j,t}^S) > \theta (C_t^S - C_t^B),$$

$$(1 - \theta) \left(\frac{\left(\tilde{\mathcal{N}}_{j,t-1}^S \right)^{\beta[-a_{j,t}]} - \left(\tilde{\mathcal{N}}_{j,t-1}^B \right)^{\beta[a_{j,t}]}}{\left(\tilde{\mathcal{N}}_{j,t-1}^S \right)^{\beta[-a_{j,t}]} + \left(\tilde{\mathcal{N}}_{j,t-1}^B \right)^{\beta[a_{j,t}]}} \right) > \theta (C_t^S - C_t^B),$$

$$(1 - \theta) \left(\left(\tilde{\mathcal{N}}_{j,t-1}^S \right)^{\beta[-a_{j,t}]} - \left(\tilde{\mathcal{N}}_{j,t-1}^B \right)^{\beta[a_{j,t}]} \right) > \dots$$

$$\dots \theta (C_t^S - C_t^B) \left(\left(\tilde{\mathcal{N}}_{j,t-1}^S \right)^{\beta[-a_{j,t}]} + \left(\tilde{\mathcal{N}}_{j,t-1}^B \right)^{\beta[a_{j,t}]} \right),$$

where we used the definition of $\mathcal{I}_{j,t}^S$ according to (14). After manipulation we find (15) to be equivalent to

$$(1 - \theta) \left(\tilde{\mathcal{N}}_{j,t-1}^S \right)^{\beta[-a_{j,t}]} - (1 - \theta) \left(\tilde{\mathcal{N}}_{j,t-1}^B \right)^{\beta[a_{j,t}]} > \dots$$

$$\dots (2\theta C_t^S - \theta) \left(\tilde{\mathcal{N}}_{j,t-1}^S \right)^{\beta[-a_{j,t}]} - (2\theta C_t^B - \theta) \left(\tilde{\mathcal{N}}_{j,t-1}^B \right)^{\beta[a_{j,t}]},$$

$$(1 - 2\theta C_t^S) \left(\tilde{\mathcal{N}}_{j,t-1}^S \right)^{\beta[-a_{j,t}]} > (1 - 2\theta C_t^B) \left(\tilde{\mathcal{N}}_{j,t-1}^B \right)^{\beta[a_{j,t}]},$$

which is identical to condition (9) in the main text.

References

- Avery, C. and P. Zemsky (1998). Multidimensional uncertainty and herd behavior in financial markets. *American Economic Review* 88(4), pp. 724–748.
- Bak, P. and M. Paczuski (1995). Complexity, contingency, and criticality. *Proceedings of the National Academy of Sciences* 92(15), 6689–6696.
- Banerjee, A. V. (1992, August). A simple model of herd behavior. *Quarterly Journal of Economics* 107(3), 797–817.

- Berg, S. (1993). Condorcet's jury theorem, dependency among jurors. *Social Choice and Welfare* 10(1), 87–95.
- Bernheim, B. D. (1994). A theory of conformity. *Journal of Political Economy* 102(5), pp. 841–877.
- Bikhchandani, S., D. Hirshleifer, and I. Welch (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy* 100(5), pp. 992–1026.
- Bikhchandani, S. and S. Sharma (2001). Herd behavior in financial markets. *IMF Staff Papers* 47, 279–310.
- Blanchard, O. J., D. Romer, M. Spence, and J. E. Stiglitz (Eds.) (2012). *In the Wake of the Crisis: Leading Economists Reassess Economic Policy*. MIT Press.
- Boland, P. J. (1989). Majority systems and the condorcet jury theorem. *Journal of the Royal Statistical Society. Series D (The Statistician)* 38(3), 181–189.
- Branch, W. and B. McGough (2011, June). Business cycle amplification with heterogeneous expectations. *Economic Theory* 47(2), 395–421.
- Brock, W., J. Lakonishok, and B. LeBaron (1992). Simple technical trading rules and the stochastic properties of stock returns. *The Journal of Finance* 47(5), 1731–1764.
- Brock, W. A. and C. H. Hommes (1997). A rational route to randomness. *Econometrica* 65(5), 1059–1095.
- Brock, W. A. and C. H. Hommes (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22(8-9), 1235–1274.
- Brunnermeier, M. K. (2009). Deciphering the liquidity and credit crunch 2007-2008. *Journal of Economic Perspectives* 23(1), 77–100.
- Chiarella, C., R. Dieci, and L. Gardini (2006). Asset price and wealth dynamics in a financial market with heterogeneous agents. *Journal of Economic Dynamics and Control* 30(9-10), 1755–1786.
- Cipriani, M. and A. Guarino (2014). Estimating a structural model of herd behavior in financial markets. *American Economic Review* 104(1), 224–51.
- Devenow, A. and I. Welch (1996). Rational herding in financial economics. *European Economic Review* 40(3-5), 603 – 615. Papers and Proceedings of the Tenth Annual Congress of the European Economic Association.

- Edey, M. (2009). The global financial crisis and its effects. *Economic Papers: A journal of applied economics and policy* 28(3), 186–195.
- Epstein, J. M. (2007). *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton University Press.
- Evstigneev, I. V., T. Hens, and K. R. Schenk-Hoppe (2009). Evolutionary finance. In T. H. Klaus Reiner Schenk-Hoppe (Ed.), *Handbook of Financial Markets: Dynamics and Evolution*, Handbooks in Finance, pp. 507 – 566. San Diego: North-Holland.
- Eyster, E. and M. Rabin (2010). Naive herding in rich-information settings. *American Economic Journal: Microeconomics* 2(4), 221–43.
- Eyster, E. and M. Rabin (2014). Extensive imitation is irrational and harmful. *Quarterly Journal of Economics* 129(4), 1861–1898.
- Fama, E. F. (1965). The behavior of stock-market prices. *Journal of Business* 38(1), pp. 34–105.
- Farmer, J. D. and S. Skouras (2013). An ecological perspective on the future of computer trading. *Quantitative Finance* 13(3), 325–346.
- Foucault, T. (2012). *Algorithmic Trading: Issues and Preliminary Evidence*, pp. 1–40. John Wiley & Sons Ltd.
- Gabaix, X. and R. Ibragimov (2011). Rank - 1 / 2: A simple way to improve the ols estimation of tail exponents. *Journal of Business & Economic Statistics* 29(1), 24–39.
- Galimberti, J. K. and M. L. Moura (2016). Improving the reliability of real-time output gap estimates using survey forecasts. *International Journal of Forecasting* 32(2), 358–373.
- Gopikrishnan, P., V. Plerou, L. A. N. Amaral, M. Meyer, and H. E. Stanley (1999). Scaling of the distribution of fluctuations of financial market indices. *Physical Review E: Statistical, Nonlinear, and Soft Matter Physics* 60(5), 5305–5316.
- Grazzini, J. and M. Richiardi (2015). Estimation of ergodic agent-based models by simulated minimum distance. *Journal of Economic Dynamics and Control* 51(0), 148 – 165.
- Grofman, B., G. Owen, and S. L. Feld (1983). Thirteen theorems in search of the truth. *Theory and Decision* 15(3), 261–278.
- Hodrick, R. J. and E. C. Prescott (1997). Postwar U.S. business cycles: An empirical investigation. *Journal of Money, Credit and Banking* 29(1), 1–16.

- IMF (2015, October). Global financial stability report - vulnerabilities, legacies, and policy challenges: Risks rotating to emerging markets. Technical report, International Monetary Fund, Washington.
- Janis, I. L. (1982). *Groupthink: psychological studies of policy decisions and fiascoes*. Cengage Learning.
- King, R. G. and S. T. Rebelo (1993). Low frequency filtering and real business cycles. *Journal of Economic Dynamics and Control* 17(1-2), 207 – 231.
- Kirilenko, A. A. and A. W. Lo (2013). Moore’s law versus Murphy’s law: Algorithmic trading and its discontents. *Journal of Economic Perspectives* 27(2), 51–72.
- Kirman, A. (1993, February). Ants, rationality, and recruitment. *Quarterly Journal of Economics* 108(1), 137–156.
- Kurz, M. (1994). On the structure and diversity of rational beliefs. *Economic Theory* 4(6), 877–900.
- Kurz, M., H. Jin, and M. Motolese (2005). The role of expectations in economic fluctuations and the efficacy of monetary policy. *Journal of Economic Dynamics and Control* 29(11), 2017 – 2065.
- Ladha, K. K. (1992). The condorcet jury theorem, free speech, and correlated votes. *American Journal of Political Science* 36(3), pp. 617–634.
- Ladha, K. K. (1995). Information pooling through majority-rule voting: Condorcet’s jury theorem with correlated votes. *Journal of Economic Behavior & Organization* 26(3), 353 – 372.
- Landemore, H. and J. Elster (Eds.) (2012). *Collective Wisdom: Principles and Mechanisms*. Cambridge University Press.
- List, C., C. Elsholtz, and T. D. Seeley (2009, March). Independence and interdependence in collective decision making: An agent-based model of nest-site choice by honeybee swarms. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364(1518), 755–762.
- Lux, T. (1995, July). Herd behaviour, bubbles and crashes. *Economic Journal* 105(431), 881–96.
- Lux, T. (2009). Stochastic behavioral asset-pricing models and the stylized facts. In T. Hens and K. R. Schenk-Hoppé (Eds.), *Handbook of Financial Markets: Dynamics and Evolution*, Handbooks in Finance, pp. 161 – 215. San Diego: North-Holland.

- Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business* 36(4), pp. 394–419.
- Park, A. and H. Sabourian (2011). Herding and contrarian behavior in financial markets. *Econometrica* 79(4), 973–1026.
- Plerou, V., P. Gopikrishnan, X. Gabaix, and H. E. Stanley (2002). Quantifying stock-price response to demand fluctuations. *Phys. Rev. E* 66, 027104.
- Scheinkman, J. A. and M. Woodford (1994). Self-organized criticality and economic fluctuations. *American Economic Review* 84(2), pp. 417–421.
- Scholz, F. W. and M. A. Stephens (1987). K-sample Anderson-Darling tests. *Journal of the American Statistical Association* 82(399), 918–924.
- Simon, H. A. (1982). *Models of Bounded Rationality*. Cambridge, MA: The MIT Press.
- Sornette, D. and W.-X. Zhou (2006). Importance of positive feedbacks and overconfidence in a self-fulfilling Ising model of financial markets. *Physica A: Statistical Mechanics and its Applications* 370(2), 704 – 726.
- Suhadolnik, N., J. Galimberti, and S. Da Silva (2010). Robot traders can prevent extreme events in complex stock markets. *Physica A: Statistical Mechanics and its Applications* 389(22), 5182–5192.
- Sumpter, D. J. and S. C. Pratt (2009). Quorum responses and consensus decision making. *Philosophical Transactions of the Royal Society B: Biological Sciences* 364(1518), 743–753.
- Surowiecki, J. (2005). *The wisdom of the crowds*. Anchor.
- Tesfatsion, L. and K. L. Judd (Eds.) (2006). *Handbook of Computational Economics II: Agent-Based Computational Economics*. North-Holland.
- Wolfram, S. (1983, Jul). Statistical mechanics of cellular automata. *Rev. Mod. Phys.* 55, 601–644.
- Young, H. P. (1988). Condorcet’s theory of voting. *American Political Science Review* 82(4), 1231–1244.