



Munich Personal RePEc Archive

A Note on institutional hierarchy and volatility in financial markets

Alfarano, Simone and Milakovic, Mishael and Raddant,
Matthias

Universitat Jaume I de Castellò

2011

Online at <https://mpra.ub.uni-muenchen.de/30902/>
MPRA Paper No. 30902, posted 16 May 2011 15:02 UTC

A Note on Institutional Hierarchy and Volatility in Financial Markets*

S. Alfarano^a, M. Milaković^{†b}, and M. Raddant^c

^aDepartment of Economics, University Jaume I, Castellón, Spain

^bDepartment of Economics, University of Bamberg, Germany

^cInstitute for the World Economy (IfW), Kiel, Germany

May 10, 2011

Abstract

From a statistical point of view, the prevalence of non-Gaussian distributions in financial returns and their volatilities shows that the Central Limit Theorem (CLT) often does not apply in financial markets. In this paper we take the position that the independence assumption of the CLT is violated by herding tendencies among market participants, and investigate whether a generic probabilistic herding model can reproduce non-Gaussian statistics in systems with a large number of agents. It is well-known that the presence of a herding mechanism in the model is not sufficient for non-Gaussian properties, which crucially depend on the details of the communication network among agents. The main contribution of this paper is to show that certain hierarchical networks, which portray the institutional structure of fund investment, warrant non-Gaussian properties for any system size and even lead to an increase in system-wide volatility. Viewed from this perspective, the mere existence of financial institutions with socially interacting managers contributes considerably to financial volatility.

Keywords: Herding, financial volatility, networks, core-periphery.

JEL codes: G10, C46, D85, E19.

*We gratefully acknowledge financial support by the Volkswagen Foundation through their grant on “Complex Networks As Interdisciplinary Phenomena.” MM is thankful for financial support from the *Research Promotion Plan 2010 of Universitat Jaume I*, received during his visit to Castellón in that year. Comments by Albrecht Irle, Jonas Kauschke, Thomas Lux, and two anonymous referees are greatly appreciated.

[†]Corresponding author: mishael.milakovic@uni-bamberg.de, +49 951 863 2582.

1 Introduction

Financial time series exhibit ubiquitous non-Gaussian statistical regularities across different countries, assets, and time frequencies. The two most prominent features concern the fluctuations in the prices of financial assets, which exhibit heavy tails and clustered volatility (see, e.g., Cont, 2001; Pagan, 1996). From a statistical point of view, the prevalence of non-Gaussian distributions in returns and their volatilities testifies to the importance of long-range correlations, which ultimately prevent the application of the Central Limit Theorem (CLT). Traditional finance has paid little, if any, attention to the origins of these statistical regularities and to the possibly most challenging question implied by the violation of the CLT: how does a complex system like the financial market actually allow for a large scale coordination of the trading positions among millions of agents? The established literature on informational cascades (see, e.g., Banerjee, 1992; Bikhchandani et al., 1992; Chamley, 2004) does not address this question because it considers a static, sequential Bayesian updating approach with a constant ‘true’ state of the world and lacks any connection to the stylized facts of financial returns. The three major strands of the agent-based finance literature, on the other hand, argue in unison that it is precisely the perpetually alternating coordination of trading strategies over time that is responsible for the stylized facts of financial returns. Yet each of the approaches has to deal with its own set of problems.

Percolation models of herd behavior exploit the properties of well known critical systems from the statistical physics literature (see, e.g., Cont and Bouchaud, 2000; Iori, 2002; Bornholdt, 2001; Stauffer and Sornette, 1999) but rely on carefully adjusted model parameters near criticality to produce non-Gaussian statistics, entirely leaving open how or why a financial market composed of millions of agents could self-organize into (and remain in) such a critical state. The second strand of models follows the seminal work of Brock and Hommes (1997) where agents interact globally rather than locally, namely through the price system and public information about the performance of strategies that is subject to noise (see, e.g., Hommes, 2006; Chang, 2007). The drawback of this class of models is that they need a careful fine-tuning of their ‘signal-to-noise ratio’ around unity in order to resemble the stylized facts. Finally the third strand, and starting point of

the present paper, is inspired by entomological experiments concerning ants' foraging behavior that Kirman (1991, 1993) utilized to propose a stochastic herding model of opinion formation among financial investors. These models endogenously create swings and herding behavior in aggregate expectations through *social agent interaction*, while the stationary distribution of the stochastic process of opinion formation describes the *statistical equilibrium* of the model.

The 'ant model' has been reasonably successful in replicating the statistical features of financial returns, but Alfarano et al. (2008) have shown analytically that Kirman's original model suffers from the problem of *self-averaging* or *N-dependence*:¹ the model's ability to replicate the stylized facts vanishes for a given parametrization when the system size N increases, a quite common feature in agent-based models that has received relatively minor attention so far (see, e.g., Aoki, 2008; Egenter et al., 1999; Lux and Schornstein, 2005). Alfarano and Milaković (2009) establish a direct link between N -(in)dependence and the communication network among agents in a generalized version of Kirman's original model. They show that the model is immune to self-averaging if the relative communication range of agents remains unchanged under an enlargement of system size. Interestingly and rather counter-intuitively, other network features like the functional form of the degree distribution, the average clustering coefficient, the graph diameter, or the extent of assortative mixing have no impact on the N -dependence property. Put differently, the average number of neighbors per agent has to increase linearly with the total number of agents N in order to overcome the problem of self-averaging in the generic herding model. Among prototypical network structures such as regular lattices, small-world, or scale-free networks (see, e.g., Newman, 2003), it is only the random graph with constant linking probability that exhibits this feature, yet random graphs are hardly ever a realistic representation of socio-economic communication networks.

After all, the results of Alfarano et al. (2008) and Alfarano and Milaković (2009) establish the model's behavior when the number of agents tends to infinity, at the same time illustrating that simple proto-typical network structures (with the exception of the empirically unsatisfactory random graph) cannot overcome the problem of N -dependence. The present paper builds

¹Aoki utilizes the terms (*non*) *self-averaging* in lieu of N -(in)dependence, and we will subsequently use both terms interchangeably.

on these insights and investigates whether a certain class of core-periphery networks might be capable of overcoming the self-averaging property of the original model. Here we consider a central network with bi-directional links between *core agents* or *opinion leaders* on one hand, and a relatively large number of uni-directionally linked *followers* in the periphery on the other. We vary the number of followers per core agent by randomly drawing from various distributions, and study the aggregate behavior of system-wide opinion dynamics under an increasing dispersion in the number of followers. In essence the hierarchical network corresponds to a weighted version of the original model. As we argue below, the weighted version is a reasonable first approximation of the institutional structure of financial fund investment. The central idea is that many investors effectively transfer control over investment decisions to fund managers who in turn are socially interacting, with the opinions of some fund managers carrying greater weight than others, for instance because they manage larger funds or have performed more successfully in the past. It turns out that the analytical mean-field prediction used in Alfarano and Milaković (2009) now significantly underestimates the volatility in system-wide opinion dynamics. The key implication of this result is that behavioral heterogeneity among interacting agents is not, as previously thought, the exclusive source of endogenously arising volatility in agent-based herding models, but that the hierarchical structure of fund investment is an important auxiliary source of financial volatility.

We take the position that investing in the presence of (actively managed) financial funds basically corresponds to the hierarchical core-periphery networks we study here. Investors who are not wealthy enough to afford a broadly diversified portfolio of assets, those who participate in retirement plans, or those who simply feel that they lack the skills or time to make investment decisions often invest in some type or other of managed fund. Effectively such agents, who correspond to followers in the periphery of the network, transfer their wealth to the fund managers in the core, and ultimately allow those to make decisions for them. If fund managers socially interact with their peers, and empirical evidence by Hong et al. (2005) and Wermers (1999) strongly suggests that this is indeed the case, we arrive at the core-periphery networks that we study in this paper.

Essentially, core-periphery networks will lead to an increase of system-wide volatility because fluctuations in a disproportionately small but central

part of the network are amplified on a system-wide level. Therefore it seems rather ironic that investors who want to ‘play it safe’ by investing in a variety of managed funds will actually end up increasing system-wide volatility if they delegate investment decisions to herding-prone fund managers.

2 Generic Herding Model

In a prototypical interaction-based herding model of the Kirman type, the agent population of size N is divided into two groups, say, X and Y of sizes n and $N - n$, respectively. The time evolution of the group sizes is modeled as a Markov chain, characterized by a pair of transition rates that are sometimes also referred to as birth and death rates. Depending on the particular financial market framework, the two groups are typically labeled as fundamentalists and chartists, or optimists and pessimists, or buyers and sellers. The basic idea is that agents change state for personal reasons or under the influence of the *neighbors* with whom they socially interact during a given time period. The transition rate for an agent i to switch from state X to state Y in the Markov chain is

$$\omega_i^- \equiv \rho_i(X \rightarrow Y) = a_i + \lambda_i \sum_{j \neq i} D_Y(i, j), \quad (1)$$

where a_i governs the possibility of self-conversion due to idiosyncratic factors, e.g. the arrival of new information, while λ_i governs the interaction strength between i and its neighbors. The function $D_Y(i, j)$ is an indicator function serving to count the number of i ’s neighbors that are in state Y ,

$$D_Y(i, j) = \begin{cases} 1 & \text{if } j \text{ is a } Y\text{-neighbor of } i, \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

hence the sum captures the (equally weighted) influence of the neighbors on agent i . Symmetrically, the transition rates in the opposite direction are given by

$$\omega_i^+ \equiv \rho_i(Y \rightarrow X) = a_i + \lambda_i \sum_{j \neq i} D_X(i, j). \quad (3)$$

Let $a = \sum_i a_i / N$ and $\lambda = \sum_i \lambda_i / N$ denote the averages of the behavioral parameters over agents, and let D denote the average number of neighbors per agent. If all links are bi-directional, $\lambda_i > 0 \forall i \in \{1, \dots, N\}$, a mean-field

argument (see Alfarano and Milaković, 2009) shows that the transition rates for a single switch on the aggregated system-wide level are

$$\omega^- = n \left(a + \frac{\lambda D}{N} (N - n) \right), \quad (4)$$

for a switch from X to Y , and symmetrically

$$\omega^+ = (N - n) \left(a + \frac{\lambda D}{N} n \right), \quad (5)$$

for the reverse switch. An important result of the mean-field approach is that the *relative communication range* D/N ultimately determines whether the Markov chain is self-averaging or not. In the jargon of Alfarano et al. (2008), the non self-averaging case corresponds to “non-extensive” transition rates with a constant relative communication range, while the “extensive” transition rates, as in Kirman’s original model, lead to self-averaging and hence to counter-factual statistics of returns.² Notice that non-extensive transition rates depend on the respective *occupation numbers* n and $N - n$, while extensive transition rates depend on the *concentrations* n/N and $(N - n)/N$ of agents in the opposite state, and therefore on the average communication range per time period in the network. This apparently minor modification has a crucial impact on the aggregate properties of the herding model, as illustrated in Figure 1. Hence, in contrast to Kirman’s original model, the generalized transition rates (4) and (5) illustrate that network structure matters because the average number of neighbors explicitly enters the transition rates.

At any time, the *state of the system* refers to the concentration of agents in one of the two states, say, $z = n/N$, which can be treated as a continuous variable for large N . None of the possible states of $z \in [0, 1]$ is an equilibrium in itself nor are there multiple equilibria in the orthodox economic sense. Equilibrium rather refers to the stationary distribution of the birth and death process (4) and (5). The distribution, that is the statistical equilibrium, describes the proportion of time the system spends in state z and is known to be a Beta distribution (see, e.g., Alfarano et al., 2008; Alfarano

²The next section explains in more detail how the Markov chain typically enters Walrasian models of the financial market.

and Milaković, 2009, for a detailed derivation of the following results),

$$p_\epsilon(z) = \frac{1}{B(\epsilon, \epsilon)} z^{\epsilon-1} (1-z)^{\epsilon-1}, \quad (6)$$

where $B(\epsilon, \epsilon) = \Gamma(\epsilon)^2/\Gamma(2\epsilon)$ is Euler's Beta function. The qualitative behavior of the process is parsimoniously encoded in the adimensional shape parameter ϵ of the distribution

$$\epsilon = \frac{aN}{\lambda D}. \quad (7)$$

When $\epsilon < 1$, the distribution is bimodal with probability mass having maxima at $z = 0$ and $z = 1$. For $\epsilon > 1$ the distribution is unimodal, and in the “knife-edge” scenario $\epsilon = 1$ the distribution is uniform. The mean $E[z] = 1/2$ is independent of ϵ but the system exhibits very different characteristics depending on the modality of the distribution. In the bimodal case, the system spends least of its time around the mean, instead mostly exhibiting very pronounced herding in either of the extreme states, as illustrated in the top panel of Figure 1. Finally, the variance of z ,

$$Var(z) = E(z^2) - E(z)^2 = \frac{1}{4(2\epsilon + 1)} = \left[4 \left(\frac{2aN}{\lambda D} + 1 \right) \right]^{-1}, \quad (8)$$

is known to be a convenient summary measure of the model properties with respect to an enlargement of system size. If the variance of z remains constant (or even increases) when the system is enlarged, the leptokurtosis and volatility clustering of returns will be preserved in a standard Walrasian model of market clearing. A decreasing variance under enlargement of system size, on the other hand, is characteristic of self-averaging and thus leads to counter-factual Gaussian properties of returns, as shown in the bottom panel of Figure 1 and explained in more detail in the following section.

3 Financial Market Framework

For the sake of completeness, we briefly discuss how the Markov chain of the previous section would enter into a parsimonious model of an artificial financial market with interacting heterogeneous agents, where it is typically used as a metaphor of information diffusion among investors (see, e.g., Kirman,




Figure 1: The two panels on the top illustrate the time evolution of aggregate opinion dynamics measured as the fraction of agents in one of the two states, say, $z = n/N$ (top panel $N = 500$, $a = 0.5$, $\lambda = 1$, $D = N$; second panel $N = 10000$, $a = 0.5$, $\lambda = 1$, $D = 500$). The two panels on the bottom exhibit the corresponding time series of returns generated from a Walrasian pricing function, as for instance in Eq. (13) of Section 3 (with $\kappa = 1$), where the level of excess demand depends on z . The bottom panel illustrates that an enlargement of system size under extensive transition rates will lead to counter-factual Gaussian returns and absence of volatility clustering.

1991, 1993; Alfarano et al., 2005; Alfarano and Lux, 2007; Alfarano et al., 2008; Alfi et al., 2009; Irle et al., 2011, for more realistic or detailed implementations). Suppose that market participants are divided into two groups: the first group is populated by N_F *fundamentalists*, who buy (sell) assets when the price is below (above) its fundamental value P_F . Their excess demand for assets is given by $ED_F = N_F \gamma_F \log(P_F/P)$, where $\gamma_F > 0$ des-

ignates the sensitivity to deviations between the fundamental value and the market price P . Without loss of generality, the fundamental price is assumed to be constant over time. The second group is populated by N_{NT} *noise traders*, who are essentially driven by herd instincts in their investment strategies. Depending on their expectations of future price movements, noise traders can be either *optimists* (*subscript O*) or *pessimists* (*subscript P*). The excess demand of the noise trader group will be proportional to their aggregated state, $ED_{NT} = \gamma_{NT}(N_O - N_P)$, where N_O and N_P are the numbers of optimists and pessimists, respectively, with $N_{NT} = N_O + N_P$. The parameter $\gamma_{NT} > 0$ governs the impact of the noise traders' aggregate mood on the asset price. In line with the notation of the previous section, ED_{NT} can be parameterized as a function of $z = N_O/N_{NT}$, that is the fraction of optimists over the total number of noise traders

$$ED_{NT} = \gamma_{NT} \cdot N_{NT}(2z - 1). \quad (9)$$

While the share of fundamentalists and noise traders is constant over time (so there are no transitions between those two groups), switches from optimism to pessimism and vice versa do take place among the noise traders, and are governed by the Markov chain detailed in Section 2. Hence noise traders change their opinions about the future prospects of an asset for idiosyncratic reasons or because of a tendency to follow the majority opinion of their peers.

Assuming sluggish price adjustments by a market maker in the presence of excess demand, one typically formalizes the price dynamics as

$$\frac{dP}{P \cdot dt} = \theta \cdot ED = \theta[ED_F + ED_{NT}], \quad (10)$$

where θ is the speed of price adjustment. As an approximation to the resulting disequilibrium dynamics, one may consider instantaneous market clearing ($\theta \rightarrow \infty$) or equivalently a Walrasian scenario ($ED = 0$) and solve (10) for the equilibrium price

$$P = P_F \exp \left[\frac{N_{NT} \cdot \gamma_{NT}}{N_F \cdot \gamma_F} (2z - 1) \right] = P_F \exp [\kappa(2z - 1)], \quad (11)$$

where

$$\kappa = \frac{N_{NT} \cdot \gamma_{NT}}{N_F \cdot \gamma_F}. \quad (12)$$

Given a realization of the process z , we can see from (11) that periods of undervaluation (compared to the fundamental price) will alternate with episodes of overvaluation. In the first case the majority of noise traders are pessimists, while in the second case most are optimists.

Finally, returns are typically defined as the log-increment of prices

$$r(t, \Delta t) = \log \left(\frac{P(t + \Delta t)}{P(t)} \right) = \kappa \Delta z, \quad (13)$$

and the third panel of Figure 1 shows the corresponding time series of log-returns for a ‘small’ number of traders ($N_F = N_{NT} = 500$), visually already indicating a leptokurtic return distribution and volatility clustering. In fact, Alfarano and Lux (2007) have shown that this very simple model quantitatively reproduces the stylized facts of financial returns with (i) a fat-tailed distribution of returns, (ii) an absence of auto-correlation in raw returns, and (iii) a slowly decaying positive auto-correlation in even functions of returns, i.e. volatility clustering. Increasing the number of agents, for instance to $N_F = N_{NT} = 10,000$ (keeping $D = 500$) as shown in the bottom panel of Figure 1, turns the abrupt mood swings in z into much smoother paths and results in counter-factual Gaussian fluctuations of returns.

4 Network Hierarchy and Core Weights

Essentially, we know that the relative communication range D/N in the transition rates (4) and (5) determines whether or not the model is self-averaging. Alfarano and Milaković consider prototypical networks with bi-directional links, in particular regular lattices, random graphs, small-world networks of the Watts and Strogatz (1998) type, and the scale-free networks of Barabási and Albert (1999). Among these it is merely the random graph that exhibits a constant relative communication range since in that case $D = N \ell$, where ℓ designates the constant linking probability among agents in the random graph. On the other hand, D/N approaches zero for an increasing system size in the other network structures, unless one appropriately changes the respective parameters in the generating mechanisms of these networks.

From a socio-economic viewpoint, however, it is not at all clear how or why a complex system composed of many interacting agents could possibly coordinate an appropriate system-wide change in these parameters. The




Figure 2: A stylized representation of a hierarchical core-periphery network, where core agents (black; bi-directional links) influence each other in their opinion formation, while peripheral followers (grey; uni-directional links) simply mimic their respective core agents. This basically corresponds to a weighted version of Kirman’s original model.

random graph is not a convincing mapping of socio-economic relationships either, because it implies that the average connectivity of agents increases linearly with system size.³ Now suppose instead that N core agents are still bi-directionally linked among themselves, i.e. they still obey the Markov chain in Section 2. Additionally, each core agent has a constant number W of followers in the periphery, with uni-directional links emanating from the core to the periphery. Uni-directional linking implies that the state of peripheral followers corresponds to the state of the respective core agents. Then the total number of followers is $W N = F$, with a total of $F + N$ agents in the entire network. In this case, the system-wide concentration of agents in state X will be

$$z = \frac{Wn + n}{F + N} = \frac{n(W + 1)}{N(W + 1)} = \frac{n}{N}, \quad (14)$$

which just amounts to a relabeling of variables. Put differently, in this special case the system size $F + N$ can be expanded at will by simply adding followers F without encountering the self-averaging problem. Thus we preserve system-wide fluctuations in a population of $F + N$ individuals, although only N socially interacting core agents are responsible for the fluctuations. At

³A simple example illustrates this implausibility. Suppose you live in Smallville, where you closely interact with, say, thirty people. Moving to Metropolis, with a population about three hundred times the size of Smallville, a constant linking probability would imply that you now closely interact with a number of agents on the order of 10^5 .

the same time, the hierarchical structure avoids the empirically unsatisfactory random graph structure in the entire population that would otherwise be necessary to preserve non self-averaging fluctuations. The assumption of a constant number of followers per core agent, however, is quite artificial and unsatisfactory. Therefore we want to investigate more general core-periphery structures by randomly drawing the number of followers per core agent from various distributions, keeping the total number of followers constant. We examine whether or how the dynamics of z change when the dispersion of followers increases. Notice that the respective numbers of followers now act as *weights* in the opinion formation process of core agents, otherwise we recover the unweighted and already well-understood cases resulting in the large- N limit of the generalized transition rates (4) and (5). Put differently, we would like to avoid the problem of self-averaging when enlarging the system, but without taking recourse to random graphs. Therefore we turn to core-periphery networks as a stylized representation of the institutional structure of financial markets, and investigate whether these hierarchical networks overcome the problem of self-averaging when the core remains small relative to the periphery.

Figure 2 provides a stylized representation of the resulting core-periphery networks that reflect the organizational structure of managed fund investment. On one hand, peripheral agents who invest in managed funds effectively delegate all subsequent investment decisions to fund managers until they decide to withdraw their capital again. On the other hand, the fund managers in the core influence each other and are prone to herding effects, as in the empirical findings of Hong et al. (2005) or Wermers (1999). We can also interpret the number of followers per core agent as the size distribution of funds, thereby implicitly assuming that the influence of fund managers on each other's opinion is proportional to the size of the fund they are managing. While we are not aware of evidence that directly supports this assumption, the empirical size distribution of funds does in fact exhibit wide dispersion and leptokurtosis (see, e.g., Gabaix et al., 2006; Schwarzkopf and Farmer, 2008).

5 Simulation Setup and Results

It is important to recall that we can study the (non) self-averaging property without actually increasing the number of agents in our subsequent simulations, because the addition of followers amounts to changing core agent *weights*. Notice that adding core agents instead of followers would correspond to the scenario that Alfarano and Milaković (2009) already studied in detail, where the structure of the bi-directional (core) network determines whether the model is self-averaging or not in the large N limit. The introduction of weights, however, prevents a straightforward application of their mean-field technique: when the weights are widely dispersed, the average number of followers per core agent obviously no longer provides a good approximation. Therefore we simulate the opinion dynamics in various core-periphery models, where we increase the dispersion of weights while drawing weights from different distributions, or altering the network structure in the core. We compare the resulting variance of z both to the mean-field prediction and to the variance in another limiting case that we have termed the *independent one-leader scenario* below.⁴ After all, the variance of z is a useful summary measure of the different scenarios because we know that if it decreases relative to the mean-field benchmark, the weighted core-periphery networks will still suffer from the problem of self-averaging. If on the other hand the variance of z remains constant, the hierarchical model will be immune to self-averaging.

5.1 Network-adapted transition rates

To implement individual transition probabilities, in line with the transition rates (1) and (3), we first consider the (symmetric) adjacency matrix $E = e_{ij}$ for $i, j \in \{1, \dots, N\}$ that keeps track of the links or *edges* between core agents, with $e_{ij} = 1$ if i and j are neighbors and $e_{ij} = 0$ otherwise.⁵ The key element in the implementation of the transition rates is to determine for each agent i the number of neighbors that are in the opposite state, say n_i .

Let $\mathbf{e}(i)$ denote the i -th column of the adjacency matrix E , which basically informs us of who is or is not an i -neighbor. While some of the neighbors will be in the same state as agent i , others will be in the opposite

⁴Appendix A contains an analytical treatment of this case.

⁵By convention, $e_{ii} = 0$, so there are no ‘self-loops’ in the network.

state, and these are the agents that we are interested in when implementing transition rates. To extract the i -neighbors that are in the opposite state, consider the projection matrix $S(i)$ of dimension $N \times N$ that keeps track of the i -neighbors that are in the opposite state: that is, for each i the off-diagonal elements of $S(i)$ are zero, $s_{ij} = 0$ if $i \neq j$, and obviously $s_{ii} = 0$ as well; on the diagonal of $S(i)$ we have $s_{jj} = 1$ if the state of neighbor j is opposite to that of agent i , and $s_{jj} = 0$ otherwise. Then the column vector $\mathbf{k}(i) = S(i) \mathbf{e}(i)$ expresses which i -neighbors are in the opposite state, and we finally have $n_i = \mathbf{k}^T(i) \mathbf{k}(i)$.

Thus *in the absence of followers* we would posit the transition probability $\tilde{\pi}_i = (a + \lambda n_i) \Delta t$ for switching states on the individual level. To ensure that $0 \leq \tilde{\pi}_i \leq 1 \forall i$, we need $\Delta t \leq 1/(a + \lambda n_{max})$, where n_{max} designates the number of neighbors of the node(s) with the highest degree in the network. Since an agent can be connected at most to all other agents, we utilize the transition probability

$$\tilde{\pi}_i = \frac{a + \lambda n_i}{a + \lambda N} \quad (15)$$

for individual switches, hence agent i 's probability to remain in the current state is $0 \leq 1 - \tilde{\pi}_i \leq 1$.

In the presence of followers, we first need to make sure that our simulation results are comparable with the mean-field prediction arising from (15), hence the individual transition probabilities need to be adapted to the core weights stemming from the hierarchical network setup. Let the column vector \mathbf{w} , with elements w_i , record the number of followers or weights for each core agent i , so $F = \sum_i w_i$ is the total number of followers in the network, and let $\langle f \rangle = F/N$ be the average number of followers per core agent. Now we are interested in the weighted sum of core agents who are in the opposite state of an agent i , denoted f_i . Since $\mathbf{k}(i)$ describes the i -neighbors that are in the opposite state, the weighted sum of core agents in the opposite state is straightforwardly computed as $f_i = \mathbf{k}^T(i) \mathbf{w} = \mathbf{e}^T(i) S(i) \mathbf{w}$, and the probability p_i to observe a change in the state of agent i in the weighted scenario is now given by

$$\pi_i = \frac{a + \lambda f_i / \langle f \rangle}{a + \lambda N} \quad (16)$$

Notice several points about the formulation of the herding term in the numerator of (16). First, using the definition of $\langle f \rangle$, we can rewrite it as $f_i / \langle f \rangle = N f_i / F$. Since $0 \leq f_i / F \leq 1$, we see that the denominator in (16)

ensures $0 \leq \pi_i \leq 1 \forall i$. Put differently, since $0 \leq n_i \leq N$, the weighted formulation has the same boundaries as the unweighted one. Second, if all core agents have the same number of followers, we have that $\forall i f_i = n_i \langle f \rangle$, so we recover the unweighted original formulation (15). Third, the ratio $f_i / \langle f \rangle$ in the sum of (16) is a direct measure of the dispersion of core weights, readily illustrating why we should not expect the mean-field approximation to be accurate when the dispersion becomes large.

5.2 Simulation setup

In our simulations, we fix the number of core agents at $N = 500$ and draw the number of followers from Gaussian, uniform, exponential and Pareto distributions with mean $\langle f \rangle = 1000$ such that each randomly drawn number is rounded to the nearest (absolute) integer value. Let N^+ and F^+ respectively denote the number of core agents and followers that are in the optimistic state. The system-wide concentration of agents in the optimistic state is now $z = (N^+ + F^+) / M$, where $M = N + F$ is the total number of agents. In all scenarios we set the parameters a, λ in such a way that $\epsilon = 1$, which yields a uniform distribution of z with $Var(z) = 1/12 \approx .083$ when the mean-field approximation applies. One ‘sweep’ of the system corresponds to one round of sequential updating of all core agents in the system, thus requiring N steps per sweep, and each simulation run consists of half a million sweeps. Finally, we successively increase the standard deviation σ_f of the respective distribution while ensuring that the weights remain positive and record the variance of z for each sequence of increasing σ_f . Recall again that when $Var(z)$ increases (decreases) above (below) the “knife-edge” value of one twelfth, this implies that the distribution of z transforms from a uniform to a bimodal distribution with non-trivial averaging behavior (unimodal distribution with trivial self-averaging).

5.3 Core structure and one-leader benchmark

The simulation results for a fully connected core are magnified in the inlay of Figure 3. When core weights are not overly dispersed, the mean-field prediction still performs well, but pronounced deviations ultimately do occur as the dispersion of weights increases. Intuitively, this happens because a few core agents become increasingly influential in the opinion formation




Figure 3: The impact of increasing heterogeneity in core weights on system-wide volatility. The simulations (with a fully connected core $D = N$ and behavioral parameters $a = \lambda = 1$, so $\epsilon = 1$) demonstrate that rising heterogeneity leads to increasing volatility, irrespective of the particular distribution from which the weights are drawn.

dynamics of the system, thereby increasing the time during which the system is near one of the two extreme states. Hence hierarchical networks are not only immune to self-averaging, but actually *amplify* volatility in the system. It is noteworthy that the outcome does not depend on the functional form of the distribution from which the weights are drawn.

In order to determine the limit of the variance amplification, we consider an extreme case that we label as the *one-leader* scenario. In this case, we allocate an equal number of followers to all but one core agent (the leader), who is then assigned a weight such that the average number of followers corresponds again to $\langle f \rangle = 1000$. Let $1/N < p < 1$ denote the fraction of followers that are connected to the one-leader, such that the leader has pF followers, and assume that the remaining $(1 - p)F$ followers are allocated with equal weight among the $N - 1$ remaining core agents. When $p = 1/N$, all core agents have the same number of followers, F/N . Conversely when $p \rightarrow 1$, the system is almost entirely represented by the leader. In each

new simulation run, we successively shift a larger number of followers to the leader by increasing p . The result is shown in Figure 3, with $Var(z)$ intuitively approaching a value of one-fourth since the leading agent will represent almost the entire system by itself, and cannot be influenced by others anymore. Thus its actions will consist of random switches between the two states, while the few remaining core agents mimic the leader’s behavior. Hence the system spends half its time in one state and half in the other, resulting in a variance of one-fourth.

In addition, we present some analytical results for the related benchmark scenario of an *independent* leader who in a sense acts “outside” the core network: in contrast to the preceding one-leader scenario, the independent leader now does not care about the state of other core agents. We explain the details of this benchmark setup and how we simulated it in Appendix A. Figure 3 also illustrates the benchmark outcome of both the prediction and the simulation for the independent leader scenario. In summary, the figure establishes two central results. First, the mean-field approximation works reasonably well if the dispersion in the number of followers is relatively small, that is when the core is fairly homogeneously weighted. Second, a heterogeneously weighted core actually leads to increasing system-wide volatility, asymptotically approaching the independent one-leader scenario, which constitutes the most extreme degree of heterogeneity. The heterogeneity of core agents thus represents an auxiliary source of fluctuations in the model. Basically, the simulations establish the model’s behavior between the mean-field and independent leader benchmarks, illustrating that the core-periphery setup quickly diverges from the mean-field approximation and asymptotically approaches the independent leader benchmark. The social interactions among core agents are crucial for overcoming N -dependence, because a vanishing herding propensity would lead to independently acting core agents, and thereby to a degenerate self-averaging distribution of their aggregate opinion with a sharply peaked mean of one-half. Moreover, if the core was enlarged (instead of the periphery) we would also confront the self-averaging problem, unless the core network was the empirically unsatisfactory random graph. It is therefore the contemporaneous presence of a hierarchical network with a relatively small core, and the social interactions in the core that ultimately overcomes the problem of N -dependence.

5.4 Varying the core structure

Our previous investigations show that a hierarchical network with a fully connected core not only overcomes the self-averaging problem, but also amplifies volatility. A remaining issue is whether these results are robust with respect to the network structure in the core itself. On that account, we perform another series of simulations with varying core network structures, and record how the different core networks respond to an increasing dispersion of weights.

For comparison with our previous findings we keep the size of the core fixed at $N = 500$, and construct the following networks in the core: a circle with neighborhood forty, a random network with linking probability of ten percent, and a scale-free network with an average of five thousand links. For the random and the scale-free graph we construct ten different realizations of the core network, and run the simulations again for half a million sweeps, subsequently averaging over the ten respective core realizations. The details of the respective network parametrizations are not crucial, because in each scenario we set $\lambda = 1$ in the transition rates (16), and adapt the behavioral parameter a in light of a particular parametrization of D such that the mean-field prediction would again yield a uniform distribution ($\epsilon = 1$). The simulation results in Figure 4 demonstrate that core network structure has merely second-order effects on the macroscopic properties of the model. As before, an increasing dispersion of followers increases volatility, while the mean-field prediction holds true if the dispersion of weights is not too large.

We also simulated a very extreme scenario by considering a scale-free graph with deterministically assigned core weights that are proportional to the degree of a core agent. We can think of such a *proportional weights* structure as the asymptotic limit of *positive feedback* effects in the time evolution of the network, for instance if highly central core agents increasingly attract the interest and wealth of investors, or if core agents with a large weight become increasingly connected among their peers in the core. Whatever the ultimate reason might be for observing such a double-weighted hierarchy, it is noteworthy that volatility increases considerably compared to the other scenarios shown in Figure 4, even for very small deviations in the number of followers.

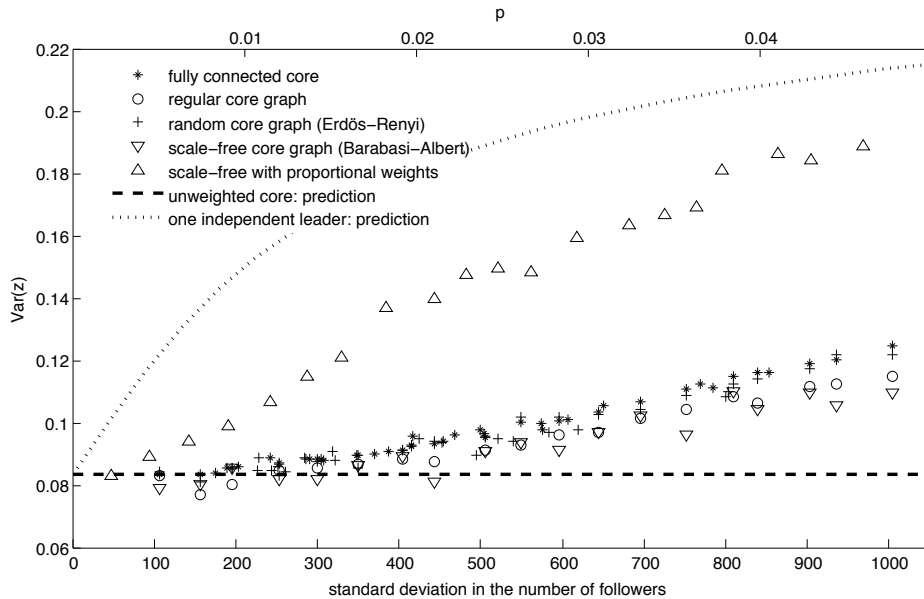


Figure 4: The impact of different core network structures on system-wide volatility has merely second-order effects, except for the “proportional weights” scenario: if the core is a scale-free graph with core weights that are proportional to the degree of each agent in the core, the mean-field approximation immediately fails to produce accurate results and the variance of z rapidly increases, almost doubling compared to the other scenarios where core weights are still randomly assigned when varying the network structure in the core. As before, the simulations were conducted for $N = 500$ agents with $\lambda = 1$, and a and D set in such a way that $\epsilon = 1$.

6 Conclusions

Hierarchical core-periphery structures turn out to overcome the problem of N -dependence in probabilistic herding models of the Kirman type. On one hand, this is good news from the viewpoint of the model’s asymptotic properties, because one is able to replicate the stylized facts of financial returns with behaviorally heterogeneous agents for any system size, without having to tune any of the behavioral parameters. On the other hand, our findings have somewhat stark implications from the viewpoint of investment strategy, and they also raise pressing new questions about the origins of hierarchical network structures.

The introduction of hierarchical network structures represents an additional source of volatility on top of the behavioral heterogeneity that has

previously been considered as the exclusive source of volatility in social interaction models. If one accepts our premise that hierarchical networks are a useful representation of fund investor relationships in financial markets, then popular and traditional investment advice to ‘diversify one’s portfolio’ has to be judged with caution. Investors who are not wealthy enough to broadly diversify their portfolios, those who participate in funded retirement plans, or those who simply feel that they lack the skills or time to make appropriate investment decisions might very well delegate their investment decisions to institutional investors. But if these fund managers are socially interacting and influencing each other in their investment decisions, as the quoted empirical evidence suggests, this becomes a self-defeating strategy because we have argued that system-wide volatility increases under such circumstances. Put in more provocative terms, all the good intentions of investors to diversify risk can lead to the opposite effect if fund managers are prone to social interaction effects. Moreover, the presence of positive feedback effects in the time evolution of hierarchical networks seems to worsen the situation further, rather than improving it, since positive feedbacks would significantly increase the level of volatility in our simulations.

From the viewpoint of policy-making, our study indicates that a reduction of financial volatility would be facilitated by a shrinking degree of hierarchical organization in financial markets, corresponding to an increasing decentralization of investment decisions. While such advice sounds straightforward in principle, its implementation would most likely be more painful and complex: our results suggest that already very small values of p (or market share for that matter) lead to a sudden and pronounced increase in volatility. Keeping p very close to zero, on the other hand, would more or less imply getting rid of managed funds altogether, which hardly appears to be a feasible option.

References

- S. Alfarano and T. Lux. A noise trader model as a generator of apparent financial power laws and long memory. *Macroeconomic Dynamics*, 11(S1): 80–101, 2007.
- S. Alfarano and M. Milaković. Network structure and N -dependence in

- agent-based herding models. *Journal of Economic Dynamics and Control*, 33:78–92, 2009.
- S. Alfarano, T. Lux, and F. Wagner. Estimation of agent-based models: The case of an asymmetric herding model. *Computational Economics*, 26:19–49, 2005.
- S. Alfarano, T. Lux, and F. Wagner. Time-variation of higher moments in a financial market with heterogeneous agents: An analytical approach. *Journal of Economic Dynamics and Control*, 32:101–136, 2008.
- V. Alfì, M. Cristelli, L. Pietronero, and A. Zaccaria. Minimal agent based model for financial markets i. *European Physical Journal B*, 67:385–397, 2009.
- M. Aoki. Thermodynamic limits of macroeconomic or financial models: One- and two-parameter Poisson-Dirichlet models. *Journal of Economic Dynamics and Control*, 32:66–84, 2008.
- A. V. Banerjee. A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3):797–817, 1992.
- A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.
- S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, 1992.
- S. Bornholdt. Expectation bubbles in a spin model of markets: Intermittency from frustration across scales. *International Journal of Modern Physics C*, 12:667–674, 2001.
- W. A. Brock and C. H. Hommes. A rational route to randomness. *Econometrica*, 65:1059–1095, 1997.
- C. P. Chamley. *Rational Herds*. Cambridge University Press, New York, 2004.
- S.-K. Chang. A simple asset pricing model with social interactions and heterogeneous beliefs. *Journal of Economic Dynamics and Control*, 31:1300–1325, 2007.

- R. Cont. Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001.
- R. Cont and J. P. Bouchaud. Herd behaviour and aggregate fluctuations in financial markets. *Macroeconomic Dynamics*, 4:170–196, 2000.
- E. Egener, T. Lux, and D. Stauffer. Finite-size effects in Monte Carlo simulations of two stock market models. *Physica A*, 268:250–256, 1999.
- X. Gabaix, P. Gopikrishnan, V. Plerou, and H. E. Stanley. A theory of power-law distributions in financial market fluctuations. *Quarterly Journal of Economics*, 121(2):461–504, 2006.
- U. Garibaldi, M. A. Penco, and P. Viarengo. An exact physical approach to market participation models. In R. Cowan and N. Jonard, editors, *Heterogeneous Agents, Interactions and Economic Performances*, Lecture Notes in Economics and Mathematical Systems, pages 91–103. Springer, Berlin, 2003.
- C. H. Hommes. Heterogeneous agent models in economics and finance. In L. Tesfatsion and K. L. Judd, editors, *Handbook of Computational Economics*, volume 2, pages 1109–1186. North-Holland, Amsterdam, 2006.
- H. Hong, J. D. Kubik, and J. C. Stein. Thy neighbor’s portfolio: Word-of-mouth effects in the holdings and trades of money managers. *Journal of Finance*, 60:2801–2824, 2005.
- G. Iori. A micro-simulation of traders’ activity in the stock market: the role of heterogeneity, agents’ interactions and trade friction. *Journal of Economic Behavior and Organization*, 49:269–285, 2002.
- A. Irle, J. Kauschke, T. Lux, and M. Milaković. Switching rates and the asymptotic behavior of herding models. *Advances in Complex Systems*, forthcoming, 2011.
- A. Kirman. Epidemics of opinion and speculative bubbles in financial markets. In M. P. Taylor, editor, *Money and Financial Markets*, pages 354–368. Blackwell, Cambridge, 1991.
- A. Kirman. Ants, rationality, and recruitment. *Quarterly Journal of Economics*, 108:137–156, 1993.

- T. Lux and S. Schornstein. Genetic learning as an explanation of stylized facts of foreign exchange markets. *Journal of Mathematical Economics*, 41:169–196, 2005.
- M. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.
- A. Pagan. The econometrics of financial markets. *Journal of Empirical Finance*, 3:15–102, 1996.
- Y. Schwarzkopf and J. D. Farmer. Time evolution of the mutual fund size distribution. *arXiv/0807.3800*, 2008.
- D. Stauffer and D. Sornette. Self-organized percolation model for stock market fluctuation. *Physica A*, 271:496–506, 1999.
- D. J. Watts and S. H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393:440–442, 1998.
- R. Wermers. Mutual fund herding and the impact on stock prices. *Journal of Finance*, 54(2):581–622, 1999.

A Independent One-Leader Benchmark

Let us start by considering an arbitrary agent in the fully connected core who is always in a *fixed* state and does not change opinion. As before, let $1/N < p < 1$ denote the fraction of followers that are connected to the fixed-opinion agent, or *independent leader*, such that the agent has pF followers, and assume that the remaining $(1-p)F$ followers are allocated with equal weight among the remaining core agents, indexed by $i = 1, \dots, N - 1$. When $p = 1/N$, all core agents have the same number of followers, F/N . Conversely when $p \rightarrow 1$, the system is almost entirely represented by the leader. For ease of notation, let us write the transition probability as

$$\pi_i = (a + \lambda N F_i / F) \Delta t \quad \text{for } i = 1, \dots, N - 1, \quad (17)$$

where F_i now denotes the *system-wide number of followers in the opposite state*, so the $N - 1$ equally weighted core agents will obey the transition rate (17). As before, we set $\Delta t = 1/(a + \lambda N)$. Furthermore, let β be an

indicator function that takes on the values 0 or 1 depending on whether the state of agent i equals or is different from the state of the fixed-opinion agent. Then we can rewrite the herding term in Eq. (17), NF_i/F , taking into account the fixed opinion of the leader (say, being optimistic)

$$N \frac{F_i}{F} = \frac{N}{F} \left(Fp\beta + (n-1) \frac{F(1-p)}{N-1} \right), \quad (18)$$

which yields the modified version of the transition probability (17),

$$\begin{aligned} \pi_i &= \left(a + \lambda Np\beta + \lambda \frac{N}{N-1} (1-p)n \right) \Delta t \\ &\approx (a + \lambda Np\beta + \lambda(1-p)n) \Delta t \end{aligned} \quad (19)$$

for large N . Depending on the value of the indicator function β , the transition probabilities of agent i are either

$$\pi_i = (\varepsilon + (1-p)n) \lambda \Delta t \quad \text{or} \quad (20)$$

$$\pi_i = (\varepsilon + Np + (1-p)n) \lambda \Delta t, \quad (21)$$

where we adapted ε to the definition (7) by noting that a fully connected core implies $D = N$.

Fixing the opinion of one agent is equivalent to creating an asymmetry in the autonomous component that stems from the additional term Np in the modified transition rates. Put simply, the system exhibits a tendency towards the fixed opinion that depends on p . A straightforward mean-field argument (see, e.g., Alfarano and Milaković, 2009) results in the following system-wide transition probabilities, analogous to an extensive version of the transition rates (4) and (5),

$$\pi^- = \frac{n}{N} \frac{\varepsilon + (1-p)(N-n)}{\varepsilon + N}, \quad (22)$$

$$\pi^+ = \frac{(N-n)}{N} \frac{\varepsilon + Np + (1-p)n}{\varepsilon + N}. \quad (23)$$

The equilibrium distribution of such a unary Markov process is (see, e.g., Garibaldi et al., 2003) the *Polya distribution* $R(\varepsilon_1, \varepsilon_2; z)$, with $z = n/N$ and

shorthands⁶

$$\varepsilon_1 = \frac{\varepsilon + Np}{1 - p}, \quad \varepsilon_2 = \frac{\varepsilon}{1 - p}. \quad (24)$$

Increasing the value of the control parameter p leads to an increasingly asymmetric distribution peaked around the opinion of the leader. Fixing the opinion of one agent, however, yields a very unsatisfactory approximation for the simulations in Section 5, where the leader is not in a fixed state but rather switches states as well. Therefore we proceed by assuming that the ‘independent’ leader switches opinion randomly, without being influenced by other agents, which basically means that the autonomous term in the mean-field transitions (22) and (23) is now stochastic and time-dependent, hinging on the random realizations of the leader’s state.

Such a situation is harder to tackle analytically because it leads to a stochastic differential equation with random coefficients. In order to approximate the full mathematical problem, we employ a so-called *adiabatic* approximation that neglects the adjustment of the system to the switching of the leader by assuming that the leader’s switches are slow enough in order for the $N - 1$ agents to reach statistical equilibrium. Then we can consider the system as being in statistical equilibrium most of the time and, consequently, the resulting equilibrium distribution G_e becomes the superposition of two independent equilibrium distributions, corresponding to the two possible configurations of the leader,

$$G_e = \frac{1}{2}R(\varepsilon_1, \varepsilon_2; z) + \frac{1}{2}R(\varepsilon_2, \varepsilon_1; z), \quad (25)$$

which is an average of the previous asymmetric distributions among the two alternative configurations of the leader. The equilibrium distribution is now symmetric (note the interchange of the parameters ε_1 and ε_2) and U-shaped. From Eq. (25), the second moment of the equilibrium distribution $M_{2,e}$ is given by

$$M_{2,e} = \frac{1}{2}M_2(\varepsilon_1, \varepsilon_2) + \frac{1}{2}M_2(\varepsilon_2, \varepsilon_1), \quad (26)$$

where $M_2(\cdot, \cdot)$ denotes the second moment of the respective asymmetric Polya distribution with parameters $\varepsilon_1, \varepsilon_2$, and the variance of the equilib-

⁶The Polya distribution converges to the Beta distribution for large N . The results of this section, however, do not significantly depend on whether we use a continuous or discrete approach (material upon request).

rium distribution for a given p is

$$\begin{aligned} \text{Var}[z]_p &= \frac{1}{2}\text{Var}[\varepsilon_1, \varepsilon_2] + \frac{1}{2}\text{Var}[\varepsilon_2, \varepsilon_1] \\ &+ \frac{1}{2} \{M_1^2[\varepsilon_1, \varepsilon_2] + M_1^2[\varepsilon_2, \varepsilon_1]\} - \left(\frac{1}{2}\right)^2, \end{aligned} \quad (27)$$

where M_1 designates the first moment of the respective asymmetric Polya distribution, and $1/2$ is obviously the mean of the equilibrium distribution G_e . The two variances are equal since they are the same under an exchange of the two parameters $\varepsilon_1, \varepsilon_2$, hence the previous equation can be written as

$$\text{Var}[z]_p = \text{Var}[\varepsilon_1, \varepsilon_2] + \frac{1}{2} \{M_1^2[\varepsilon_1, \varepsilon_2] + M_1^2[\varepsilon_2, \varepsilon_1]\} - \frac{1}{4}. \quad (28)$$

It is possible to show (see, e.g., Garibaldi et al., 2003) that

$$M_1[\varepsilon_1, \varepsilon_2] = \frac{\varepsilon_1}{\varepsilon_1 + \varepsilon_2}, \quad (29)$$

$$\text{Var}[\varepsilon_1, \varepsilon_2] = \frac{\varepsilon_1 \varepsilon_2}{(\varepsilon_1 + \varepsilon_2)^2} \frac{\varepsilon_1 \varepsilon_2 + N}{N(\varepsilon_1 \varepsilon_2 + 1)}, \quad (30)$$

and utilizing these in Eq. (28) yields

$$\text{Var}[z] = \frac{1}{4} - \frac{\varepsilon_1 \varepsilon_2}{(\varepsilon_1 + \varepsilon_2)(\varepsilon_1 + \varepsilon_2 + 1)}. \quad (31)$$

Finally, recalling the shorthands in (24), we obtain the variance as a function of the control parameter p ,

$$\text{Var}[z]_p = \frac{1}{4} - \frac{1 + pN}{(2 + pN)(3 + pN - p)}, \quad (32)$$

under the parameter choice $\varepsilon = 1$, i.e. $\lambda = 1$ and $a = 1$. For $N \gg 1$, we immediately see that Eq. (32) provides boundary values that are consistent with our previous findings: if $p = 1/N$, the variance tends to $1/12$, representing the correct value for the uniform distribution (recall the parameter choice $\varepsilon = 1$); if $p \rightarrow 1$, the variance tends to $1/4$, representing a distribution concentrated either in 0 or 1. We simulated the modified model with a randomly switching leader, successively increasing the control parameter p in a fully connected core of size $N = 500$ with a total of $F = 500,000$ followers. As before, we simulated each parametrization with half a million

sweeps.

The results, along with the prediction (32), are shown in Figure 3. For easier comparison with the simulation results in Figures 3 and 4, we calculate the standard deviation in the number of followers for each parametrization of p , which is

$$\sigma = \sqrt{\frac{1}{N}(pF)^2 + \frac{N-1}{N} \left(\frac{(1-p)F}{N-1} \right)^2 - \left(\frac{F}{N} \right)^2}, \quad (33)$$

and invert the relation to obtain

$$p = \frac{\sigma\sqrt{N-1}}{F} + \frac{1}{N}. \quad (34)$$

While the independent leader scenario exhibits a quicker convergence to the limiting variance of one-fourth than the one-leader model, both versions are qualitatively similar in the sense that there is sudden and pronounced increase in volatility already for small values of p . The main difference between the independent vs one-leader scenarios is that the independent leader switches randomly (thus independently) between the two states, while the switches of the one-leader in Section 5.3 still depend on the interactions with the other core agents, which intuitively slows down the variance amplification relative to the independent leader case. As far as the creation of fluctuations and therefore risk is concerned, the important common feature of both models is that they exhibit a sudden and pronounced increase in system-wide volatility as soon as a relatively small number of core agents obtains a disproportionately large weight in the core network.