Network Hierarchy in Kirman's Ant Model: Fund Investment Can Create Systemic Risk*

S. Alfarano[a], M. Milaković[b], and M. Raddant[b]

[a]Department of Economics, University of Castellón, Spain
[b]Department of Economics, University of Kiel, Germany

This version: January 2009

Abstract

Kirman's "ant model" has been used to characterize the expectation formation of financial investors who are prone to herding. The model's original version suffers from the problem of N-dependence: its ability to replicate the statistical features of financial returns vanishes once the system size N is increased. In a generalized version of the ant model, the network structure connecting agents turns out to determine whether or not the model is N-dependent. We investigate a class of hierarchical networks in the generalized model that presumably reflect the institutional heterogeneity of financial markets. These network structures do overcome the problem of N-dependence, but at the same time they also increase system-wide volatility. Thus network structure becomes an auxiliary source of volatility in addition to the behavioral heterogeneity of interacting agents.

*We gratefully acknowledge financial support by the Volkswagen Foundation through their grant on "Complex Networks As Interdisciplinary Phenomena." Alfarano thanks the Department of Electrical and Electronics Engineering at the University of Cagliari for its hospitality during the period Dec08–Jan09.

1 Introduction

Inspired by entomological experiments concerning ants' foraging behavior, Kirman (1991, 1993) proposed a stochastic model of opinion formation among financial investors. The model endogenously creates swings and herding behavior in aggregate expectations through *agent interaction*, while the stationary distribution of the stochastic process of opinion formation corresponds to the *statistical equilibrium* of the model. The "ant model" has been reasonably successful in replicating the statistical features of financial returns, like volatility clustering and power-law tails of the returns distribution (see, e.g., Alfarano et al., 2005), but Alfarano et al. (2008) have shown analytically that Kirman's original model suffers from the problem of N-dependence: the model's ability to replicate the stylized facts vanishes for a given parametrization once the system size N is increased, a feature that although not uncommon in agent-based models, has received relatively minor attention so far (see Aoki, 2008; Egenter et al., 1999; Lux and Schornstein, 2005; Alfi et al., 2008).[1] Alfarano and Milaković (2009) establish a direct link between the problem of N-dependence and network structure among agents in a generalized version of Kirman's ant model. They prove that the model is immune to N-dependence if the *relative communication range* of agents remains unchanged under an enlargement of system size.[2] Put differently, the average number of neighbors per agent has to increase linearly with the total number of agents N in order to overcome N-dependence. Among the standard

[1] Aoki utilizes the terms *(non) self-averaging* in lieu of N-(in)dependence.

[2] 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.

prototypical network structures (see, e.g., Newman, 2003), such as regular lattices, small-world or scale-free networks, it is only the random graph with constant linking probability that exhibits this feature, but random graphs are hardly a realistic way of describing socio-economic relationships. Alfarano and Milaković conclude that hierarchical core-periphery structures might be capable of overcoming N-dependence, having the additional advantage that they are likely to be more realistic representations of the institutional or structural relationships of socio-economic interaction.

In the present paper we build on these insights and investigate whether certain core-periphery structures are indeed immune to N-dependence. The core-periphery networks that we consider here consist of a core with bi-directional links between core agents, or *opinion leaders*, and a relatively large number of *followers* who are uni-directionally linked to core agents. We vary the number of followers per core agent by randomly drawing from various distributions, and study the aggregate behaviour of system-wide opinion dynamics under an increasing dispersion in the number of followers. It turns out that the analytical mean-field prediction used by Alfarano and Milaković, which yields accurate predictions among standard network structures, now significantly underestimates the volatility in system-wide opinion dynamics. One noteworthy implication of this result is that behavioral heterogeneity among interacting agents is not the only source of endogenously arising volatility, but that the network structure describing the very feasibility of agent interaction is another potential source of volatility as well.

It appears reasonable that hierarchical core-periphery networks resemble the institutional structures in fund investment behavior, therefore we will use the terms *hierarchy* and *institutional heterogeneity* interchangeably. 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 network, transfer their wealth to fund managers in the core, and put them in charge of subsequent investment decisions on their behalf. We end up with a stylized hierarchical network of the type described above if fund managers are indeed influencing each other in their decision making, and recent empirical evidence by Hong et al. (2005) in fact documents that fund managers who work in geographical proximity are prone to what they term "word-of-mouth" effects. Essentially core-periphery networks lead to an increase of system-wide volatility, relative to a collection of independently acting agents, because fluctuations in a disproportionately small but central part of the network are amplified on a system-wide level if the network is characterized by the structural or institutional heterogeneity outlined above. Therefore it seems rather ironic that investors who want to "play it safe" by investing in managed funds will actually contribute to systemic risk if they delegate investment decisions to socially interacting fund managers.

2 The 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. Depending on the model setup, 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 their *neighbors*, with whom they interact during a given time period. The transition rate for an agent i to switch from state X to state Y is

$$\pi_i^- \equiv \rho_i(X \to Y) = a_i + \lambda_i \sum_{j \neq i} D_Y(i,j), \tag{1}$$

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 neighbor j. The function $D_Y(i,j)$ is an indicator function serving to count the number of i's neighbors that are in state Y,

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

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

$$\pi_i^+ \equiv \rho_i(Y \to X) = a_i + \lambda_i \sum_{j \neq i} D_X(i,j) \ . \tag{3}$$

If all links are bi-directional, $b_i > 0 \, \forall i \in \{1, \dots, N\}$, Alfarano and Milaković (2009) demonstrate that the transition rates for a single switch on the system-wide level are given by

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

for a switch from X to Y, and

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

for the reverse switch, where the parameters a and b are now the ensemble averages of the corresponding individual parameters a_i and b_i, and D is the average number of neighbors per agent.

If the *relative communication range* D/N remains constant under an enlargement of system size, the model is immune to N-dependence. In the jargon of Alfarano et al. (2008), this case corresponds to "non-extensive" transition rates, while the "extensive" transition rates in Kirman's original model are N-dependent. 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. This apparently minor modification has a crucial impact on the macroscopic 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 shows up explicitly in the transition rates.

At any time, the *state of the system* refers to the concentration of agents in one of the two states, say, $z = n/N$. None of the possible states of $z \in [0, 1]$ is an equilibrium in itself,[3] nor are there multiple equilibria in the orthodox sense. Equilibrium rather refers to the stationary distribution of the process given by (4) and (5), yielding the proportion of time the system spends in

[3]Notice that for large N, the concentration can be treated as a continuous variable.

state z. This statistical equilibrium distribution turns out to be (see Alfarano and Milaković, 2009, for a detailed derivation) a Beta distribution,

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

where $B(\epsilon, \epsilon) = \Gamma(\epsilon)^2/\Gamma(2\epsilon)$ is Euler's Beta function, and the shape parameter of the distribution is given by

$$\epsilon = \frac{aN}{\lambda D}. \tag{7}$$

For $\epsilon < 1$, 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 becomes uniform. The mean value of z, $E[z] = 1/2$, 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 the time around the mean, mostly exhibiting very pronounced herding in either of the extreme states, as shown in the top panel of Figure 1. Finally, the variance of z,

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

is 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 simple Walrasian market clearing scenario, while a decrease of the variance under enlargement of system size leads to counter-factual Gaussian properties of returns, as shown in the bottom panel of Figure 1.

3 Hierarchical Network Structure

The relative communication range D/N in the transition rates (4) and (5) determines whether or not the model is N-dependent. Alfarano and Milaković (2009) 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 hard to see how or why a complex system composed of many interacting agents could possibly coordinate an appropriate system-wide change in these parameters. The 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.[4] A simple way to preserve immunity from N-dependence, without taking recourse to fully connected or random networks, is to in-

[4] A 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.

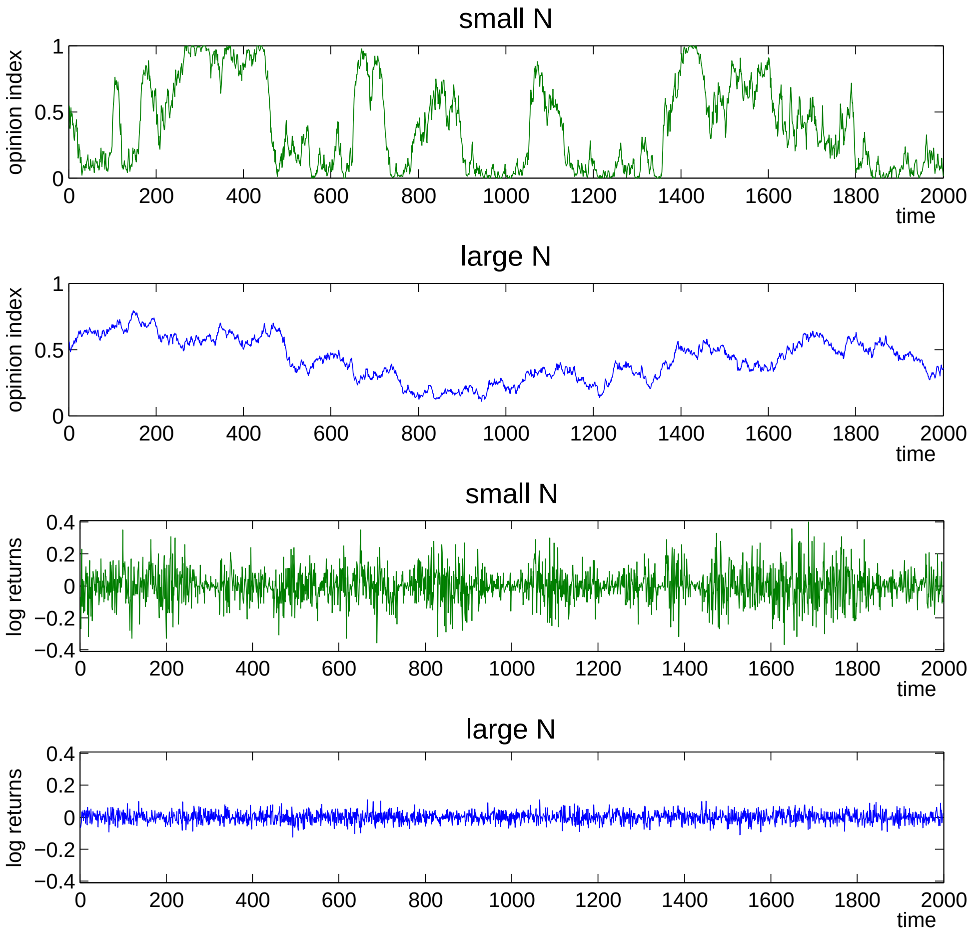

Figure 1: The two panels on the top illustrate the time evolution of aggregate opinion dynamics measured as the fraction of agents in one of the two states, say, $z = n/N$. The two panels on the bottom show the corresponding time series of (log) returns generated from a Walrasian pricing function where the level of excess demand depends on z. An enlargement of system size under extensive transition rates leads to counter-factual Gaussian returns. Non-extensive transition rates can reproduce the extensive "small N" scenario for any system size such that the pronounced swings in aggregate opinion dynamics and the resulting statistical properties of returns, like leptokurtosis and volatility clustering, are preserved under an enlargement of system size (see Alfarano et al., 2008, for more details).

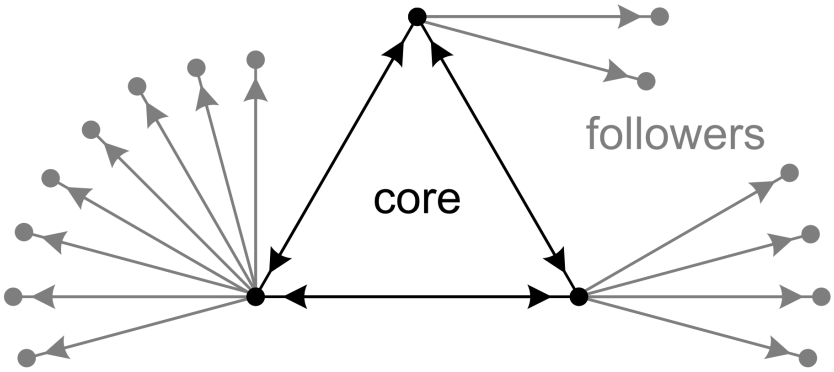

Figure 2: A stylized representation of a hierarchical core-periphery network, where core agents (black; bi-directional links) influence each other in their opinion formation, while peripheral agents (grey; uni-directional links) simply mimic their respective core agents.

troduce a hierarchical core-periphery structure. Suppose N core agents are bi-directionally linked among themselves, and each has a constant number M of followers in the periphery, with uni-directional links emanating from the core to the periphery. The uni-directional links imply that the state of peripheral agents simply corresponds to the state of their respective core agents. Then the total number of followers is $M\,N = F$, 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

$$z = \frac{Mn + n}{F + N} = \frac{n(M+1)}{N(M+1)} = \frac{n}{N}, \tag{9}$$

which amounts to a relabeling of variables. Notice that now the system size can be expanded without running into the problem of N-dependence by simply adding followers at will.

The assumption of a constant number of followers per core agent, however, is quite artificial and unsatisfactory. Therefore we investigate more general core-periphery structures by randomly drawing the number of followers from various distributions, studying whether or how the opinion dynamics change when the dispersion of followers increases. Notice that the respective numbers of followers have to act as *weights* in the opinion formation process among core agents, otherwise we would merely recover the already well-understood cases resulting from the generalized transition rates (4) and (5).

Figure 2 provides a stylized representation of the resulting core-periphery networks that presumably reflect the organizational structure of fund investment. On one hand, agents who invest in a fund effectively delegate all subsequent investment decisions to fund managers until they decide to withdraw their capital. On the other hand, fund managers in the core influence each other and are prone to herding effects, an assumption that is in fact supported by the empirical results of Hong et al. (2005) and 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 in the opinion formation dynamics is proportional to the size of the fund they are managing. While we have no direct evidence to support this assumption, the empirical size distribution of funds indeed exhibits wide dispersion and even leptokurtosis (see, e.g., Gabaix et al., 2006; Schwarzkopf and Farmer, 2008).

4 The Simulation

It is paramount to realize that we can study the N-dependence property without actually increasing the number of agents in our subsequent simulations. The simple reason being that the addition of followers corresponds to changing relative core agent *weights*.[5] The introduction of weights, however, prevents a straightforward application of analytical mean-field techniques if the weights are widely dispersed because the average number of followers per core agent no longer provides a good approximation. Therefore we simulate the opinion dynamics in the hierarchical core-periphery models with an increasing dispersion of weights, and compare the outcomes to the mean-field prediction of Alfarano and Milaković (2009). If the variance of z decreases relative to their benchmark case, hierarchical networks will still suffer from the problem of N-dependence. Conversely, if the variance increases or remains unchanged, the hierarchical model will be immune to N-dependence.

4.1 Network-adapted transition rates

At the individual level, contrary to the mesoscopic description of the system in (4) and (5), an agent either remains in its current state, or switches to the other state. To implement individual transition probabilities in the absence of followers, corresponding to the transition rates (1) and (3), Alfarano and Milaković posit the transition probability $\tilde{p_i} = (a + \lambda\, n_i(j))\Delta t$ for switching states on the individual level, where $n_i(j)$ counts the number of i-neighbors j that are in the opposite state. To ensure that all agents act on the same time scale, and also that $0 \leq \tilde{p_i} \leq 1\, \forall i$, this necessitates that

$$\Delta t \leq 1/(a + \lambda\, n_{max})\,,$$

[5]Adding core agents instead of followers corresponds 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 N-dependent or not.

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, they utilize the transition probability

$$\tilde{p}_i = \frac{a + \lambda n_i(j)}{a + \lambda N} \tag{10}$$

for individual switches, and correspondingly an agent's probability to remain in the current state is $0 \leq 1 - \tilde{p}_i \leq 1$.

In order to ensure that our simulation results are comparable with the mean-field prediction arising from (10), we adapt the individual transition probabilities so as to reflect the presence of followers in our hierarchical networks. Let f_i denote the number of followers of core agent $i \in \{1, \ldots, N\}$, where $F = \sum_i f_i$ 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. If $f_i(j)$ denotes the number of followers of an i-neighbor j, then the adapted probability p_i to observe a change in the state of agent i is now given by

$$p_i = \frac{a + \lambda \sum_{j=1}^{n_i(j)} f_i(j) / \langle f \rangle}{a + \lambda N} \tag{11}$$

Notice several points about the formulation of the sum in (11). First, using the definition of $\langle f \rangle$, we can rewrite the sum as

$$N \sum_{j=1}^{n_i(j)} f_i(j)/F \, ,$$

and realizing that $0 \leq \sum_{j=1}^{n_i(j)} f_i(j)/F \leq 1$, we see that the denominator in (11) ensures $0 \leq p_i \leq 1 \, \forall i$. Put differently, since $0 \leq n_i(j) \leq N$, the new measure should have the same boundaries, which is true for the sum in (11). Second, if core agents have the same number of followers, $\forall i \, f_i = \langle f \rangle$, we recover the original formulation (10), consistent with the result concerning the relabeling of variables. Third, the ratio $f_i(j)/\langle f \rangle$ in the sum of (11) is a measure of dispersion around the average number of followers in the network, explicitly showing that the mean-field approximation will not be very accurate if the followers are widely dispersed among core agents.

4.2 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 the absolute value of each randomly drawn number is rounded to the nearest 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 $M = F + N$ is the total number of agents. For all the different scenarios, we choose the parameters a, λ in such a way that $\epsilon = 1$, which according to the mean-field prediction of a uniform distribution of z should yield $Var(z) = 1/12 \approx .083$. One "sweep" of the system corresponds to one round of sequential updating of all 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,[6] and record the variance of z for each sequence of increasing σ_f. If $Var(z)$ increases (decreases) above (below) one twelfth, this implies that the distribution of z transforms from a uniform to a bimodal (unimodal) distribution. Recall that in the bimodal case, the system is immune to N-dependence, while the unimodal case would imply that the hierarchical network structure still suffered from N-dependence.

4.3 A fully connected core

We obtain a fairly simple representation of individual transition rates if agents in the core are fully connected among each other since the probability to observe a switch in the state of a core agent simplifies to

$$p_i = \frac{a + \lambda N F_i / F}{a + \lambda N}\,, \tag{12}$$

where F_i denotes the system-wide number of followers in the opposite state, and we can simulate the model without explicitly keeping track of the network structure. The simulation results for a fully connected core are shown in the left panel of Figure 3.

As long as the heterogeneity in the number of followers is not very extreme, the mean-field prediction still performs well, but pronounced deviations ultimately do occur as the dispersion in the number of followers rises. Intuitively, this happens because a few core agents become increasingly in-

[6] Generally, we start from distributions that are sharply peaked around $\langle f \rangle$ and increase σ_f in twenty steps. In the Gaussian and uniform cases, we simply increase the variance in constant steps. In the exponential case, we adapt the support of the distribution to account for different variances with a mean of $\langle f \rangle$, and in the Pareto case we successively decrease the tail index to generate a higher variance.

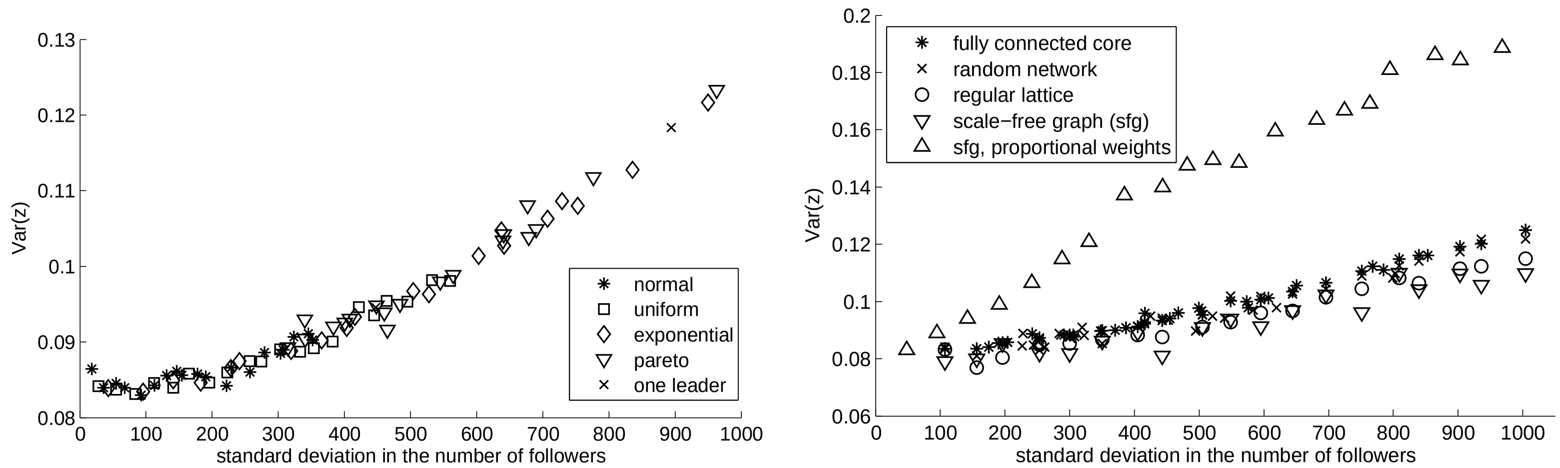

Figure 3: The impact of increasing heterogeneity in core agent weights on system-wide volatility. The simulation with a fully connected core (left panel) shows that rising heterogeneity leads to increasing volatility of z, irrespective of the particular distribution from which the weights are drawn. The right panel shows that 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 being proportional to core node degrees, the mean-field approximation immediately fails to produce accurate results, and the variance of z quickly increases, relative to the other core network structures, by a factor of up to two.

fluential in the opinion formation dynamics of the system, thereby increasing the time during which the system is near one of the two extreme states. Hence the hierarchical network structure is not only immune to N-dependence, but it actually *amplifies* volatility in the system. It is also noteworthy that the outcome does not appear to depend on the functional form of the distribution from which we randomly draw the number of followers.

In order to determine the limit of the variance amplification, we consider an extreme case that we label as the *one-leader* scenario, where 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 $\langle f \rangle = 1000$. In each new simulation run, we successively shift a larger number of followers to the leader. The result is shown in the left panel of Figure 4, with $Var(z)$ approaching a value of one-fourth, which is quite intuitive: 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, leading to a variance of one-fourth. This case also turns out to be a useful starting point for our analytical approach in Section 5.

In summary, the fully connected benchmark case establishes two central results. First, the mean-field approximation works reasonably well if the dispersion in the number of followers is relatively small. Second, a hierarchical network structure actually leads to an increase in system-wide volatility, and thereby presents an additional source of volatility in probabilistic herding models. The social interactions among core agents are crucial for overcoming N-dependence, because a vanishing herding propensity would lead to independent opinion leaders, and thereby to a degenerate distribution peaked around the mean (of one-half) that is characteristic of N-dependence. It is therefore the contemporaneous presence of hierarchy and core interactions that ultimately overcomes the problem of N-dependence.

4.4 Varying the core structure

Our previous investigations show that a hierarchical network with a fully connected core not only overcomes the problem of N-dependence, 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 adapt the behavioral parameters a and λ of the transition rates (11) such that the mean-field prediction would again yield a uniform distribution ($\epsilon = 1$). The simulation results in the right panel of Figure 3 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. The juxtaposition of the two graphs in Figure 3 readily reveals that volatility increases on the same order of magnitude as in the fully connected case.

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 evolution of the hierarchical network, for instance if highly central core agents attract the increasing interest 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 the right panel of Figure 3, even for very small deviations in the number of followers.

5 Analytical Benchmark

Naturally we are interested in deriving an analytical benchmark in order to judge the plausibility of our simulation results. Essentially, our derivation of an analytical approximation that links the degree of hierarchical organization

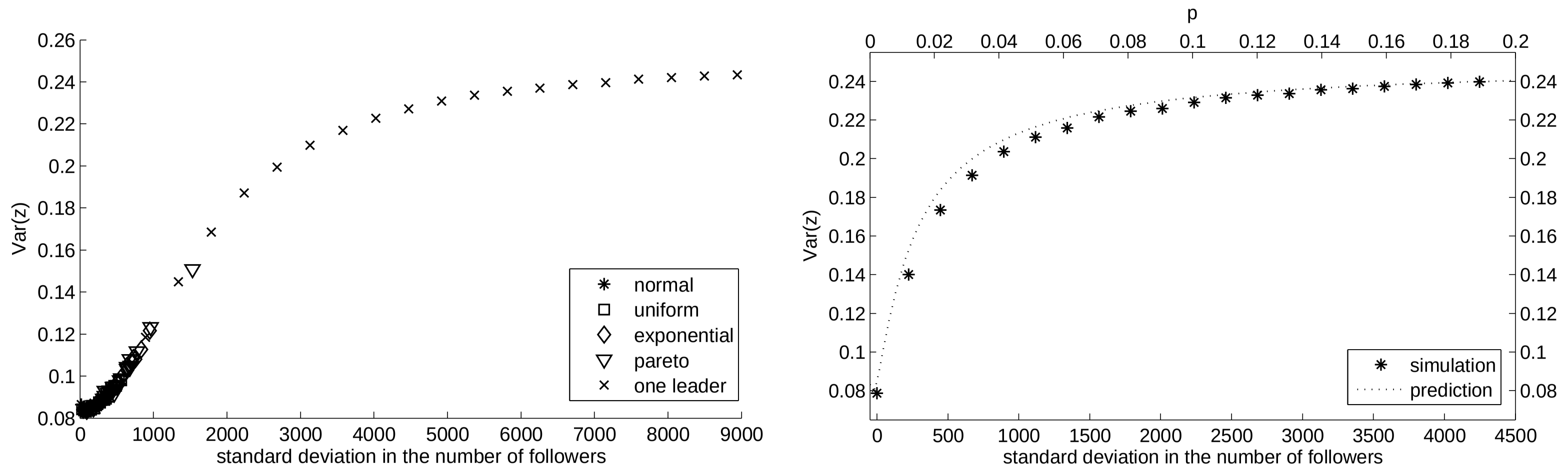

Figure 4: The left panel shows that the variance of z converges to the expected value of one-fourth in the limiting one-leader case discussed in Section 4. The right panel illustrates the analytical prediction from Eq. (27) for the p-agent scenario in Section 5 and the corresponding simulation results. For easier comparison with the left panel, we also calculated the standard deviation in the number of followers (bottom axis) for each corresponding parametrization of the control parameter p (top axis). Compared with the original setup in the left panel, the p-agent scenario displays a quicker convergence to the limiting variance of one-fourth. Nevertheless, the qualitative behavior of both model versions is very similar, exhibiting a sharp and pronounced increase in system-wide volatility when relatively few core agents have disproportionately large weights.

to systemic volatility is inspired by the one-leader benchmark in Section 4.3 because the leader can be considered as an autonomous agent whose changes of opinion are independent of the other agents. For the current purpose, let us start by considering an arbitrary agent in the core who is always in a fixed state and does not change opinion. Let $1/N < p < 1$ denote the fraction of followers that are connected to the fixed-opinion agent, or *p-agent*, such that the agent has pF followers, and assume that the remaining $(1-p)F$ followers are shared with equal weight among the remaining core agents, indexed by $i = 1, \ldots, N-1$. When $p = 1/N$, all core agents have the same number of followers, F/N. Conversely when $p \to 1$, the system is almost entirely represented by the p-agent. Notice that the $N-1$ equally weighted core agents still obey the transition rate (12).

Now let β be an indicator function that takes on the values 0 or 1 depending on whether the state of agent i is equal to or different from the state of the fixed-opinion agent. Then we can rewrite the herding term in Eq. (12), NF_i/F, taking into account the fixed opinion of the p-agent (say, optimistic)

$$N\frac{F_i}{F} = \frac{N}{F}\left(Fp\beta + (n-1)\frac{F(1-p)}{N-1}\right) \ , \tag{13}$$

which yields the modified version of the transition (12),

$$\begin{aligned} p_i &= \frac{a + \lambda N p\beta + \lambda\frac{N}{N-1}(1-p)n}{a+\lambda N} \\ &\approx \frac{a + \lambda N p\beta + \lambda(1-p)n}{a+\lambda N} \end{aligned} \tag{14}$$

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

$$p_i = \frac{\varepsilon + (1-p)n}{\varepsilon + N} \qquad \text{or} \tag{15}$$

$$p_i = \frac{\varepsilon + Np + (1-p)n}{\varepsilon + N} , \tag{16}$$

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

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 rates, analogous to an extensive version of the transitions (4) and (5),

$$\pi^- = \frac{n}{N}\,\frac{\varepsilon + (1-p)(N-n)}{\varepsilon + N} , \tag{17}$$

$$\pi^+ = \frac{(N-n)}{N}\,\frac{\varepsilon + Np + (1-p)n}{\varepsilon + N} . \tag{18}$$

The equilibrium distribution of such a unary Markov process is (see, e.g., Garibaldi et al., 2003) the *Polya distribution* $P(\varepsilon_1, \varepsilon_2; z)$, with $z = n/N$ and shorthands[7]

$$\varepsilon_1 = \frac{\varepsilon + Np}{1-p} , \qquad \varepsilon_2 = \frac{\varepsilon}{1-p} . \tag{19}$$

Increasing the value of the control parameter p leads to an increasingly asym-

[7]The 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).

metric distribution peaked around the opinion of the p-agent. Fixing the opinion of one agent, however, yields a very unsatisfactory approximation for the simulations in the previous section, where the p-agent (leader) is not in a fixed state but rather switches states as well. Therefore we proceed by assuming that the p-agent switches opinion randomly, without being influenced by other agents, which basically means that the autonomous term in the mean-field transitions (17) and (18) is now stochastic and time-dependent, hinging on the random realizations of the p-agent'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 P_e becomes the superposition of two independent equilibrium distributions, corresponding to the two possible configurations of the p-agent,

$$P_e = \frac{1}{2} P(\varepsilon_1, \varepsilon_2; z) + \frac{1}{2} P(\varepsilon_2, \varepsilon_1; z) \ , \tag{20}$$

which is an average of the previous asymmetric distributions among the two alternative configurations of the p-agent. The equilibrium distribution is now symmetric (note the interchange of the parameters ε_1 and ε_2) and U-shaped. From Eq. (20), the second moment of the equilibrium distribution $M_{2,e}$ is given by

$$M_{2,e} = \frac{1}{2}M_2(\varepsilon_1, \varepsilon_2) + \frac{1}{2}M_2(\varepsilon_2, \varepsilon_1) \ , \tag{21}$$

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 equilibrium distribution for a given p is

$$\begin{aligned} Var[z]_p &= \frac{1}{2}Var[\varepsilon_1, \varepsilon_2] + \frac{1}{2}Var[\varepsilon_2, \varepsilon_1] \\ &\quad + \frac{1}{2}\left\{M_1^2[\varepsilon_1, \varepsilon_2] + M_1^2[\varepsilon_2, \varepsilon_1]\right\} - \left(\frac{1}{2}\right)^2 \ , \end{aligned} \tag{22}$$

where M_1 designates the first moment of the respective asymmetric Polya distribution, and 1/2 is obviously the mean of the equilibrium distribution P_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

$$Var[z]_p = Var[\varepsilon_1, \varepsilon_2] + \frac{1}{2}\left\{M_1^2[\varepsilon_1, \varepsilon_2] + M_1^2[\varepsilon_2, \varepsilon_1]\right\} - \frac{1}{4} \ . \tag{23}$$

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

$$M_1[\varepsilon_1, \varepsilon_2] = \frac{\varepsilon_1}{\varepsilon_1 + \varepsilon_2} \ , \tag{24}$$

$$Var[\varepsilon_1, \varepsilon_2] = \frac{\varepsilon_1 \varepsilon_2}{(\varepsilon_1 + \varepsilon_2)^2} \frac{\varepsilon_1 \varepsilon_2 + N}{N(\varepsilon_1 \varepsilon_2 + 1)} \ , \tag{25}$$

and utilizing these in Eq. (23) yields

$$Var[z] = \frac{1}{4} - \frac{\varepsilon_1 \varepsilon_2}{(\varepsilon_1 + \varepsilon_2)(\varepsilon_1 + \varepsilon_2 + 1)} \ . \tag{26}$$

Finally, recalling the shorthands in (19), we obtain the variance as a function of the control parameter p,

$$Var[z]_p = \frac{1}{4} - \frac{1 + pN}{(2 + pN)(3 + pN - p)} \ , \tag{27}$$

under the parameter choice $\varepsilon = 1$, i.e. $\lambda = 1$ and $a = 1$. For $N \gg 1$, we immediately see that Eq. (27) 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 $p \to 1$, the variance tends to $1/4$, representing a distribution concentrated either in 0 or 1 that corresponds to the extreme one-leader case.

We simulated the modified model with a randomly switching p-agent, 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 (27), are shown in the right panel of Figure 4. For easier comparison with the original simulation results in the left panel of Figure 4, we also calculated the standard deviation in the number of followers for each parametrization of p. While the p-agent scenario exhibits a quicker convergence to the limiting variance of one-fourth than the original model, both versions are nevertheless qualitatively very similar. The main difference between the p-agent and one-leader scenarios is that the p-agent switches randomly and independently between the two states, while the switches of the one-leader in Section 4.3 still depend on the states of the other core agents, which intuitively slows down the variance amplification relative to the p-agent case. As far as the creation of systemic 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.

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 potential origins of hierarchical network structures.

The introduction of hierarchical network structures leads to an additional source of volatility, on top of the behavioral heterogeneity that has previously been considered as the exclusive source of volatility in the ant model. 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 since we have shown 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 herding. Moreover, the presence of positive feedback effects in the time evolution of hierarchical network structures seems to further worsen the situation as it significantly increases the level of volatility.

From the viewpoint of policy-making, our study indicates that a reduction of financial volatility would require a shrinking degree of hierarchical organization in financial markets, corresponding to a decentralization of investment decisions. While such advice sounds straightforward in principle, its implementation would most likely be more painful and complex because our analytical results show that values on the order of $p \approx 10^{-2}$ already 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 still hardly appears as a feasible option, notwithstanding the popularly accepted global mechanisms that led to the current financial turmoil.

References

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. Alfi, L. Pietronero, and A. Zaccaria. Minimal agent-based model for the origin and self-organization of stylized facts in financial markets. *arXiv/0807.1888*, 2008.

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.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.

E. Egenter, 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 Mathetatical Systems, pages 91–103. Springer, Berlin, 2003.

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.

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.

Y. Schwarzkopf and J. D. Farmer. Time evolution of the mutual fund size distribution. *arXiv/0807.3800*, 2008.

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.