Network Hierarchy in Kirman's Ant Model: Fund Investments Can Create Risks*

S. Alfarano M. Milaković M. Raddant

Department of Economics
University of Kiel (CAU)
D-24118 Kiel, Germany

This version: August 2008

Abstract

Kirman's "ant model" has been used to characterize the expectation formation of financial investors who are prone to herding. In Kirman's original version, however, the model suffers from the problem of N-dependence: the model's ability to replicate the statistical features of financial returns vanishes once the system size N is increased, and the network structure that describes the feasibility of agent interaction in a generalized version of the ant model determines whether or not the model suffers from N-dependence. We investigate a class of hierarchical networks and find that they do overcome the problem of N-dependence, but at the same time they also increase system-wide volatility, and therefore embody an additional source of volatility besides the behavioral heterogeneity of interacting agents. Interpreting these findings in the context of fund investment, the desire of investors to "play it safe" might increase systemic risk if core-periphery networks indeed describe the organizational structure of fund investment.

*We gratefully acknowledge financial support by the Volkswagen Foundation through their grant on "Complex Networks As Interdisciplinary Phenomena."

1 Introduction

Inspired by entomological experiments concerning ants' foraging behavior, Kirman (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, 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, e.g., Egenter et al., 1999; Lux and Schornstein, 2005; Alfi et al., 2008). Alfarano and Milaković (2008) have established a direct link between the problem of N-dependence and the network structure underlying agent interaction 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.[1] 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 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.

[1] Interestingly, and rather counter-intuitively, other network features like the functional form of the degree distribution, the average clustering coefficient, the graph diameter, the extent of assortative mixing, etc. have no impact on the N-dependence property.

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 prototypical 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 to think of hierarchical core-periphery networks as resembling the institutional structures in fund investment behavior. 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 a managed fund. Effectively such agents, who would correspond to followers in the network, transfer their wealth to fund managers (core agents), and put these managers in charge of all subsequent investment decisions until they decide to withdraw their funds again. We will 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. Since system-wide volatility increases in the presence of core-periphery networks, it seems rather ironic that investors who want to "play it safe" 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 (1993) 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 acquisition 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 (2008) 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,[2] 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 state z. This statistical equilibrium distribution turns out to be 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 $\epsilon = aN/\lambda D$ (see Alfarano and Milaković, 2008, for a detailed derivation). For $\epsilon < 1$, the distribution is bimodal, with probability mass having maxima at $z = 0$ and $z = 1$. For $\epsilon > 1$, the distri-

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

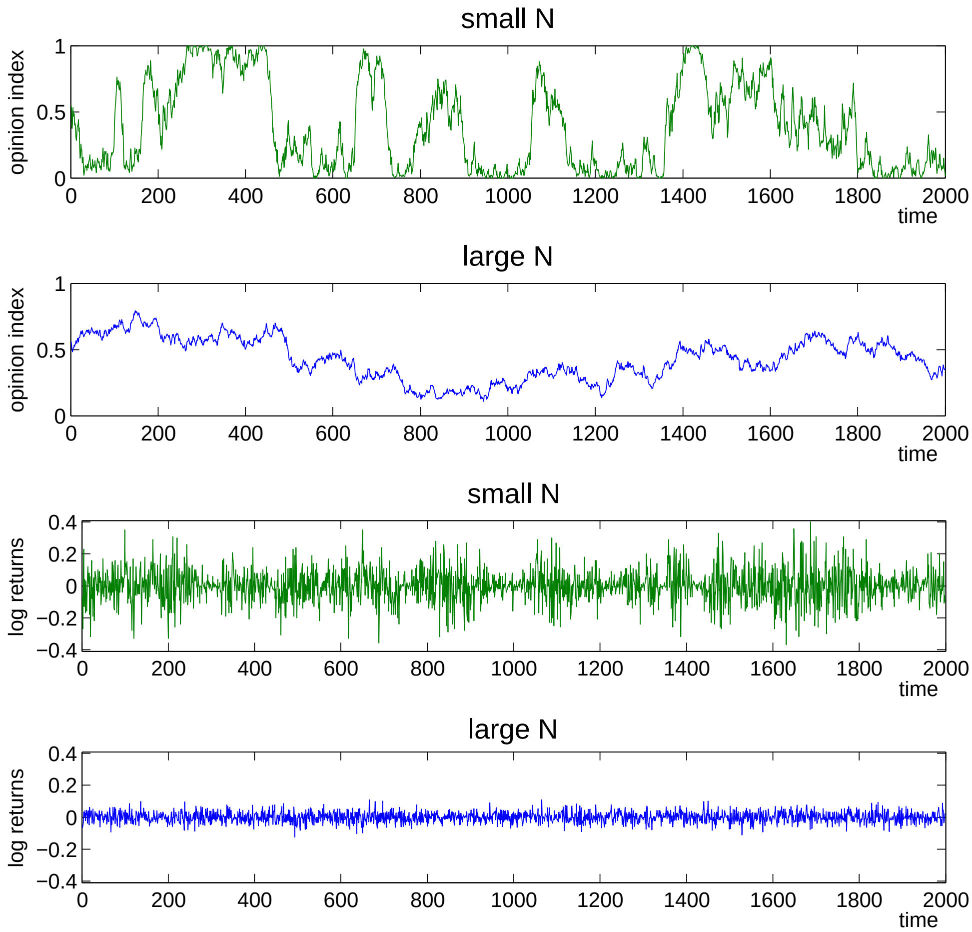

Figure 1: The two panels on the top illustrate the time evolution of aggregate opinion dynamics measured as the fraction of agents in one of the two states, say, $z = n/N$. The two panels on the bottom show the corresponding time series of (log) returns generated from a Walrasian pricing function where 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, hence 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).

bution 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,

$$Var(z) = E(z^2) - E(z)^2 = \frac{1}{4(2\epsilon + 1)} = \left[4 \left(\frac{2aN}{\lambda D} + 1 \right) \right]^{-1} , \qquad (7)$$

is a convenient summary measure of the model properties with respect to an enlargement of system size. If the variance of z remains constant 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 an enlargement of the system 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 (2008) 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 $D = N\,p$ for a random graph, where p designates the constant linking probability among agents. 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. Neither is the random graph a very convincing mapping of socio-economic relationships, because it would imply that the average connectivity of agents increases

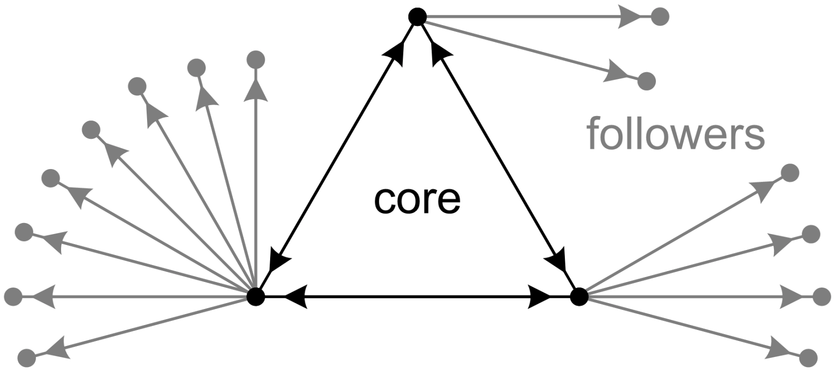

Figure 2: A stylized representation of a hierarchical core-periphery network, where core agents influence each other in their opinion formation, while peripheral agents simply mimic their respective core agents.

linearly with system size.[3] A simple way to preserve immunity from N-dependence without taking recourse to fully connected or random networks, originally suggested by Alfarano and Milaković, is to introduce 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

$$z = \frac{Mn + n}{F + N} = \frac{n(M+1)}{N(M+1)} = \frac{n}{N}, \tag{8}$$

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 therefore we want to investigate more general core-

[3]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, the random network with constant linking probability would imply that you can now closely interact with almost ten thousand people.

periphery structures by randomly drawing the number of followers from various distributions, and to study whether or how the opinion dynamics change when the dispersion of followers increases. Notice that the respective numbers of followers will 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, which might appropriately describe the organizational structure in fund investments. 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, the assumption that fund managers in the core are interacting and prone to herding effects is in fact supported by the 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, it does not appear entirely unreasonable, and there is indeed evidence that the empirical size distribution of funds exhibits wide dispersion and even leptokurtosis (see, e.g., Gabaix et al., 2006; Schwarzkopf and Farmer, 2008).

4 The Simulation

The introduction of weights 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 (2008).

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 $\Delta t \leq 1/(a + \lambda n_{max})$, 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, Alfarano and Milaković utilize the transition probability

$$\tilde{p_i} = \frac{a + \lambda n_i(j)}{a + \lambda N} \tag{9}$$

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 (9), 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

$$p_i = \frac{a + \lambda \sum_{j=1}^{n_i(j)} f_i(j) / \langle f \rangle}{a + \lambda N} \tag{10}$$

Notice several points about the formulation of the sum in (10). 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 can see that the denominator in (10) 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 (10). Second, if core agents have the same number of followers, $f_i = \langle f \rangle \; \forall i$, we recover the original formulation (9), consistent with the result concerning the relabeling of variables. Third, the ratio $f_i(j)/\langle f \rangle$ in the sum of (10) is a measure of dispersion around the average number of followers in the network, hence it shows explicitly 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^+, F^+ denote the number of core agents and followers that are, say, 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 $Var(z) = 1/12 \approx 0.83$. 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,[4] and record the variance of z for each sequence of increasing

[4]Generally, we start from distributions that are sharply peaked around $\langle f \rangle$ and increase the standard deviation in twenty steps. In the Gaussian and uniform cases, this is simply accomplished by increasing the variance in constant steps. In the exponential case we

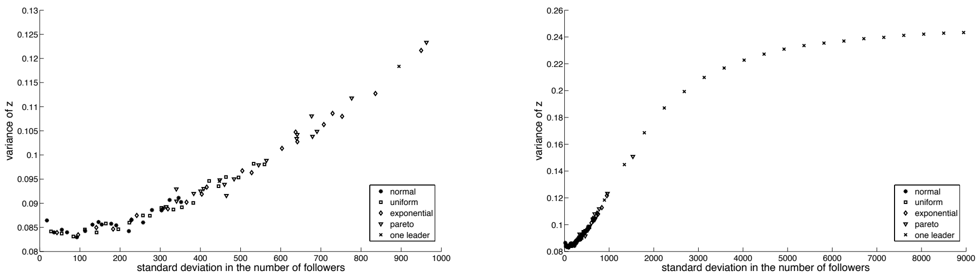

Figure 3: The impact of increasing heterogeneity in core agent weights on system-wide volatility in opinion dynamics. The left panel shows that rising heterogeneity leads to increasing volatility of z, irrespective of the particular distribution from which weights are drawn. The right panel illustrates that the variance of z converges to the expected value of one fourth in the limiting "one leader" case.

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. 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 The benchmark: 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

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

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

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 characteristic exponent to generate a higher variance.

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 influential 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 noteworthy that the outcome is not influenced by the exact functional form of the distribution from which we randomly draw the number of followers.

This allows us to determine the limit of the variance amplification by considering an extreme case that we label as the *one leader* scenario, where we assign an equal number of followers to all but one core agent, the "leader," who is in turn assigned a number of followers 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 right panel of Figure 3, with $Var(z)$ obviously approaching a value of one fourth, which is in fact what we would expect in such a case: 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 all other agents instantaneously 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.

In summary, the benchmark case establishes two central results. First, the mean-field approximation works reasonably well if the dispersion in the number of followers is not too large. 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.

4.4 Network structure in the core

Our previous investigations show that a hierarchical net with a fully connected core not only overcomes the problem of N-dependence, but also amplifies volatility. The remaining issue is whether these results are robust with

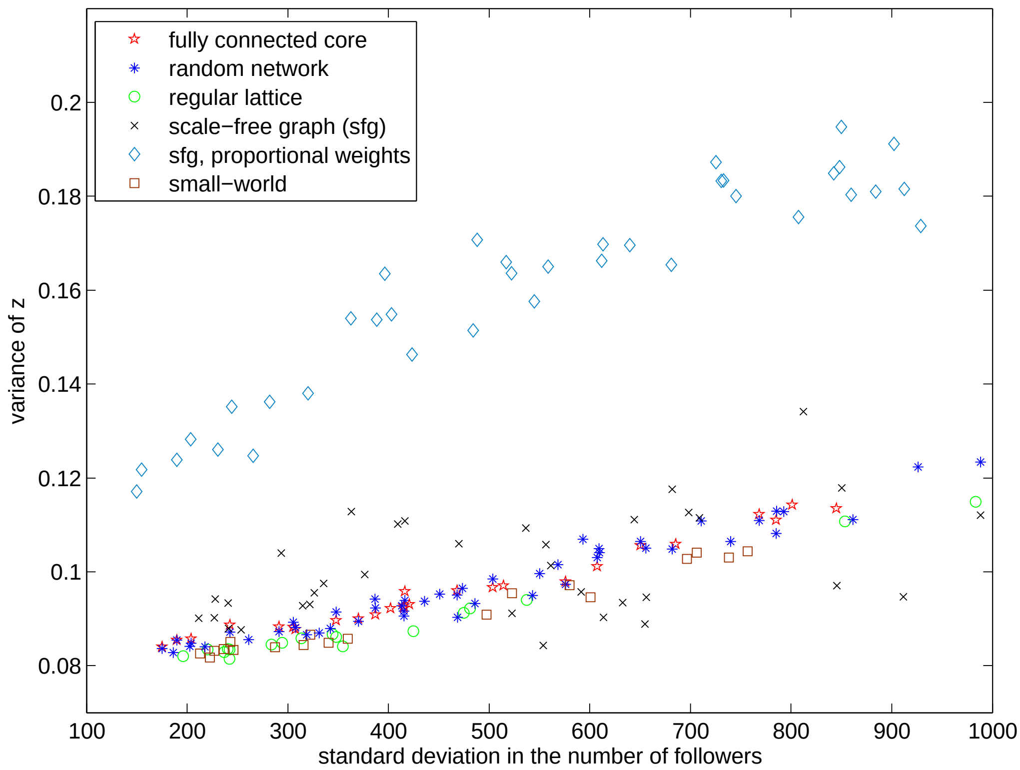

Figure 4: Impact of different network structures in the core on system-wide volatility. The "proportional weights" scenario refers to a scale-free graph in which the weights of core agents are proportional to their degree distribution in the core. In this case the mean-field approximation immediately fails to produce accurate results, and the variance of z increases almost by a factor of two.

respect to the network structure in the core itself. On that account, we perform another series of simulations, this time varying the structure in the core while again recording how the different networks respond to an increasing dispersion of weights in the core.

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 one 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 random core realizations. In each scenario, we adapt the behavioral parameters a and λ in the transition rates (10) in such a way that the mean-field prediction again yields a uniform distribution ($\epsilon = 1$). The simulation results in Figure 4 demonstrate that the network structure in the core merely has second-order effects on the macroscopic properties of the model. Once more, an increasing dispersion of followers increases volatility, while the mean-field prediction holds true if the dispersion of weights is not too large. Juxtaposing the left panel of Figure 3 with Figure 4 shows that volatility increases on the same order of magnitude as in the benchmark case.

Finally we investigate a very extreme scenario, the result of which is also exhibited in Figure 4, and labeled there as the *proportional weights* case. There we consider 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 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 rather intriguing that volatility increases quite considerably in comparison to the other scenarios that we studied.

5 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 position that hierarchical networks are a useful representation of fund investor relationships in financial markets, this would imply that 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 the investment decision to institutional investors. But if these fund managers are socially interacting and influencing each other in their investment decisions, and we cited empirical evidence that in fact supports such a scenario, then this would appear as a self-defeating strategy because we have shown that system-wide volatility increases in this case. 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 dynamical feedback effects in the time evolution of hierarchical network structures would appear to further worsen the situation since it significantly increases the level of volatility.

References

S. Alfarano and M. Milaković. Network structure and N-dependence in agent-based herding models. *Journal of Economic Dynamics and Control*, 2008. In press, doi:10.1016/j.jedc.2008.05.003.

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.

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.

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. 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.