PHYSICAL REVIEW E **77**, 011106 (2008)

Spies in the minority game

You-Yang Yu,[1] Chen Xu,[2] Guo-Qing Gu,[1] and Pak Ming Hui[3]
[1]*School of Information Science and Technology, East China Normal University, Shanghai 200062, China*
[2]*Department of Physics and Jiangsu Key Laboratory of Thin Films, Suzhou University, Suzhou 215006, China*
[3]*Department of Physics, The Chinese University of Hong Kong, Shatin, Hong Kong, China*
(Received 11 August 2007; published 8 January 2008)

We study the effects of the existence of another type of agents, called spies, in the minority game (MG). Unlike the normal agents in the MG, the spies do not carry any strategy. Instead, they decide their action by scouting some normal agents and take the minority action of the spied group. For a few spies and when there is useful information in the normal agents' actions, the spies can avoid the crowd effect of the normal agents and win more readily. When information becomes less useful and when more spies are present, the spies' crowd effect hurts the success rate of the spies themselves, and the normal agents could have a higher success rate than the spies. More spies actually assist more normal agents to win, as the spies also provide more winning quotas. This leads to a nonmonotonic behavior in the total success rate of the population as a function of the fraction of spies.

DOI: 10.1103/PhysRevE.77.011106 PACS number(s): 02.50.Le, 05.65.+b, 89.65.Gh

I. INTRODUCTION

The Minority Game (MG) is a popular model in studying adaptive complex systems [1–4] and its variations have been applied to different areas such as to simulate the behavior in financial markets [5]. It is related to the so-called El Farol bar problem [6] in which a number of people decide weekly whether to enjoy a drink in a bar or to stay home without getting into a crowd in the bar. The MG is a simpler binary version in which agents compete for a limited resource. In the original MG [1], a population of an odd number of N agents decide between two possible choices, say to attend room "0" or to attend room "1," at each round of the game. The agents belonging to the minority group, say those in room "0," are the winners in that round. The winning action (take room "0") is announced to all agents and thus serves as the common information or history bit string. The agents decide based on the most recent m winning outcomes. In each round, the agents decide according to the prediction of best performing strategy in hand. A strategy maps every possible m-bit winning outcomes to a prediction. Therefore, there are 2^{2^m} possible strategies in the whole pool of strategies or full strategy space. At the beginning of the game, each agent randomly picks S strategies from the pool, with repetition allowed. The performance of a strategy reflects how successful it is to predict the winning outcome from the beginning of the game. Note that the best performing strategy of an agent changes with time. This rather simple model shows nontrivial properties [3,7,8]. For example, the number of agents making a particular decision ("0" or "1") fluctuates in time. The standard deviation σ shows a nonmonotonic dependence on m, with a large value at small m and approaching the random coin-toss limit at large m. For intermediate value of m, there is a minimum in σ [7,9,10] signifying a collective performance of the system that is better than random. The nonmonotonic behavior of σ can be explained analytically by the crowd-anticrowd theory [11,12] and by tracing the strategy performance ranking patterns as in the strategy ranking theory [13,14].

The large value of σ at small m implies a wastage of resource and more agents could have won. Attempts have been made to improve the performance of the population. A key idea is to avoid overadaptation and thus the formation of a big crowd making the same decision. This can be achieved, for example, by having a biased pool of strategies [15] or by local-formation sharing through a network [14,16]. An alternative way is to allow for an inhomogeneous population with two or more types of agents. For example, systems with a mixed population having a different number of strategies or values of m [17] and with some agents who participate randomly into the game [18] show a higher overall success rate. The reasons are that an inhomogeneous population tends to mess up the history bit string that leads to crowd formation and having agents responding to the history differently further avoids crowd formation. Having some agents who do not use the MG strategies and instead decide by an intelligent trial and error method, Xie and Wang [19] showed that one could achieve global optimization with σ taking on the possible minimum value. Mimicking a kind of agents called contrarians in real markets, Zhong *et al.* [20] showed enhanced performance in a mixed population with normal MG agents and contrarians, who take the opposite action of what the best strategy predicts.

In real-life scenarios, it is unrealistic to have a population with all agents following the rules of a game. Some agents may be smarter or they may simply cheat. Here, we study a mixed population consisting of normal MG agents and a group of agents who do not carry any strategy. These agents, called spies, can collect information from a portion of normal agents and decide by analyzing the information. It is shown that their presence could enhance the overall success rate of the population. In Sec. II, we present our model and introduce the action of the spies. In Sec. III, we present results of numerical simulations of our model and analyze the results within a phenomenological theory. A summary of the results is given in Sec. IV.

II. MODEL

Consider an inhomogeneous population with a total of $N=N_n+N_s$ agents playing the Minority Game, with N being an odd number. In each round, the winners are those in the minority group. There are thus, in principle, a maximum of $(N-1)/2$ winners per round. There are N_n normal MG agents, i.e., they hold S strategies from the full strategy space each of which maps the most recent m-bit history into a prediction, and decide based on their temporarily best-performing strategy. They have no information on the action of the other agents except for the publicly available history bit string. There are N_s agents who do not hold any strategy. They are referred to as "spies." In each round, each of these agents is assumed to have the ability to randomly pick a group of k *normal agents* and look at what their actions are. With this information and the intention to win, a spy decides the action by taking the minority side of his observed k group of normal agents. Note that the minority side of the k group that a spy observes may not be the actual winning side. The k group for each spy, who chooses the normal agents to be spied upon entirely at random, serves to provide local information. The spies are assumed to act independently, i.e., there is no communication among the spies. The fact that real-world spies receive orders from a central agency and thus may act in a correlated way is not considered here. The performance of the system is expected to depend on the value of k and the number of spies N_s, in addition to m and S as in the basic MG. If k is too small, the information may not be useful. If N_s is large, the market impact of the spies themselves becomes important as the spies could lead to severe crowd effect.

Our model presents a different inhomogeneous population than in other models. Here, the spies can get hold of the forthcoming actions of a number of randomly selected normal agents. They do not hold any strategy and they do not adapt by learning from past experience. In other local minority games [21–23], some agents can exchange information on their past decisions with their neighbors, and an agent can follow the prediction of the neighbors with the best cumulative performance. Here, the spies can look at different groups of k normal agents in each turn and there are no fixed neighbors. Thus the model is also different from networked MGs in which an agent has a few fixed neighbors during the game. These agents thus mimic those persons in real market situations who collect, sometimes illegally, and act based on insider information.

III. RESULTS AND DISCUSSIONS

We study a system with $N_n=101$ normal agents. We define $\rho_s=N_s/N$ to be the fraction of spies and $\rho_k=k/N_n$ as the extent to which the normal agents are being spied upon for information. The normal agents carry $S=2$ strategies, which are picked at the beginning randomly from the full strategy space. A normal agent decides his action based on the prediction of the best-performing strategy that he has in hand at the moment of decision. We aim to study how the existence of spies affects the global performance of the system such as

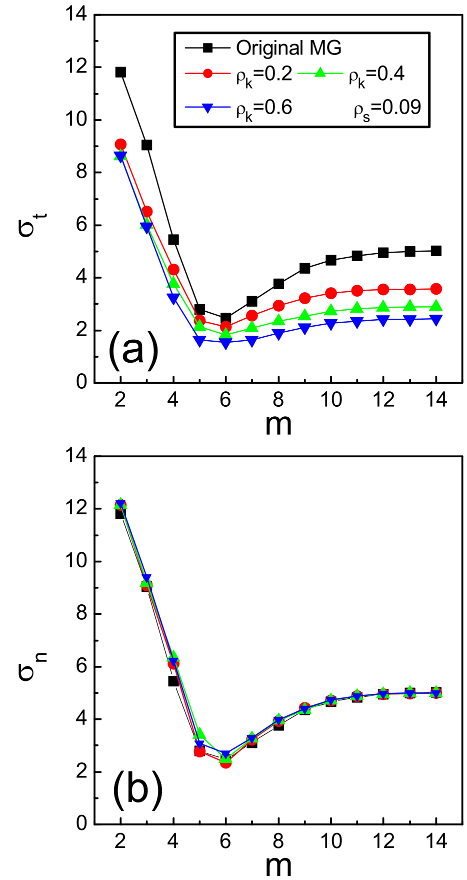

FIG. 1. (Color online) (a) The standard deviation σ_t in the number of agents making a particular decision, regardless of the type of agents, as a function of m for different values of the fraction ρ_k of normal agents from which a spy could collect information. The fraction of spies is fixed at $\rho_s=0.09$. (b) The standard deviation σ_n considering only the decisions of the normal agents, who played in a game together with the spies.

σ and the success rates. From the winning criterion of the MG, several quantities can be used to characterize the performance of the whole system. They include the standard deviation of the number of agents in the winning group, the standard deviation of the number of agents making a particular decision (either "0" or "1"), the standard deviation of the difference in the number of agents making opposite decisions, and the average winning probability. By recording the total number of agents making a particular decision, i.e., regardless of normal agents or spies, over time, the corresponding standard deviation σ_t can be calculated. Figure 1(a) shows the standard deviation σ_t as a function of m for different values of ρ_k. The number of spies is fixed at 10 and thus $\rho_s=0.09$. The behavior is typical of systems with a small fraction of spies. The simulation results are obtained by averaging over 100 independent runs with different initial distributions of strategies among the normal agents and different initial histories in starting a run. Results of the basic MG are also shown for comparison. In the basic MG, a higher σ_t implies that the agents do not make full use of the available winning quotas and a smaller σ_t implies that the number of winners is close to the allowed value. Comparing the results with that of the basic MG, the existence of a few spies leads to a lower σ_t, indicating an improved performance of the system as a whole. The increase in performance is most noticeable for small values of m. With two types of agents, it is useful to look at whether the existence of spies changes the behavior of the normal agents. Figure 1(b) shows the standard deviation σ_n in the number of normal agents making a particular decision, again for $\rho_s=0.09$. The results show that the collective behavior of the normal agents alone is largely unchanged and follows that of the basic MG with 101 agents. This indicates that a few spies do not disturb significantly the history bit string that is dominated by the actions of the normal agents. Without disturbing the behavior of the normal agents, the few spies may benefit from knowing the actions of the normal agents. The larger the k group being watched, the more useful information it is. In the MG, there is information in the history bit string and thus the actions of the agents in the small m regime. This information, however, leads to a severe crowd effect and thus a large σ_t in the basic MG. In other words, there are fewer winners than allowed and there are available winning quotas left. The spies can exploit the information and make use of these remaining winning quotas. In addition, the spies also add to the winning quotas in the game. Thus the spies' success will not hurt the normal agents' winning, leading to an overall enhancement in performance. For an intermediate value of m where σ_t in the basic MG shows a minimum, the number of normal agent winners is close to $N_n/2$ with only a few winning quotas left. The information obtained by the spies becomes less useful. The spies add some winning quotas to the game. However, these quotas can be used by both the normal agents and the spies. For large m, the situation is close to intermediate m. The information for the spies is not useful. The spies create more winning quotas that can be used by both the normal agents and the spies themselves. If the "spies" just replicate some of the agents, the results will be different. For example, if each spy simply follows a single normal agent and different spies are allowed to follow different normal agents, the crowd effect is more severe and thus σ_t is larger than the original MG. If we introduce instead a kind of agent who plays one of his strategies randomly in every turn, these agents can sometimes avoid the crowd effect, but sometimes enhance the crowd effect. The action of the spies in the present model is based on the intention to avoid the "crowd" using the limited information collected in the k group being spied upon.

As the number of spies in the population increases, the situation becomes more complicated. There are several effects. The size of the whole population is increased and hence the possible number of winners allowed by the MG also increases. The spies are making rooms for themselves and the normal agents to win. In addition, the market impact of the spies enhances, i.e., the actions of the spies become increasingly important in deciding the actual winning outcome. This in turn leads to a crowd effect due to the spies. The history bit-string of the game will then be different from that in the basic MG, due to the presence of the spies. Figure 2 shows the standard deviation σ_t as a function of m, for different values of the fraction of spies ρ_s. The spies scout $k=20$ normal agents. When there is no definite action suggested by the k group, a spy decides the action randomly. For small m, as the number of spies increases, the averaged standard deviation σ_t decreases for small fractions of spies then increases at large fractions of spies [see Fig. 2(a)]. The standard deviation σ_n deviates from that of the original MG when the fraction of spies becomes large. It is found that at $\rho_s \gtrsim 0.23$, σ_n starts to deviate from the behavior in the original MG, especially in the intermediate range of m [see Fig. 2(b)], where the low σ_n in the original MG takes on a large value.

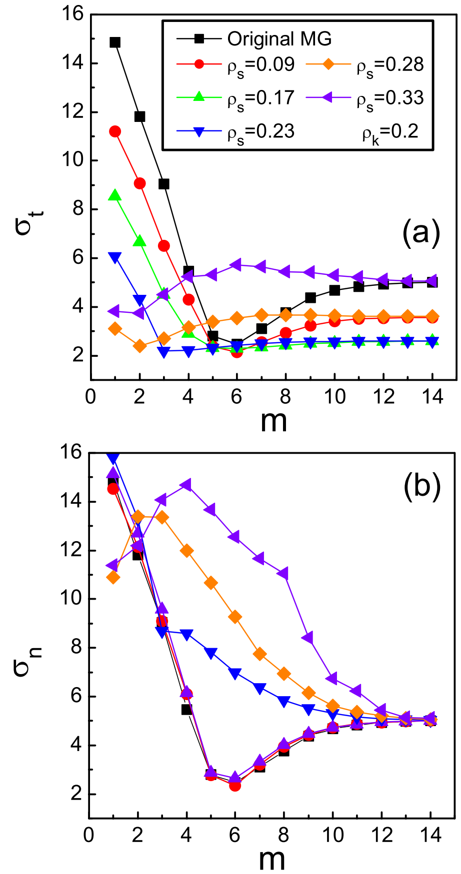

FIG. 2. (Color online) (a) The standard deviation σ_t of the whole population as a function of m, for different fractions of spies ρ_s in the population. The spies could collect information from $k=20$ normal agents ($\rho_k=0.17$). (b) The standard deviation σ_n considering only the decisions of the normal agents, who played in a game together with the spies.

Next, we study the success rate, which is the winning probability per agent per turn, in the system. It will be interesting to see how the presence of spies affects the success rates of the normal agents and the spies themselves. In Figs. 3(a)–3(c), we show the success rate of the normal agents R_n and the success rate of the spies R_s as a function of increasing fraction of spies ρ_s in the population, for three different values of $m=2$, 6, and 10 that correspond to different regimes in the basic MG. The spies can scout $k=20$ normal agents. For $m=2$ [Fig. 3(a)], the success rate R_n among the normal agents remains almost a constant, which is about 0.4 as in the basic MG, when the number of spies is small and the success rate R_s among the spies is quite high (~0.7).

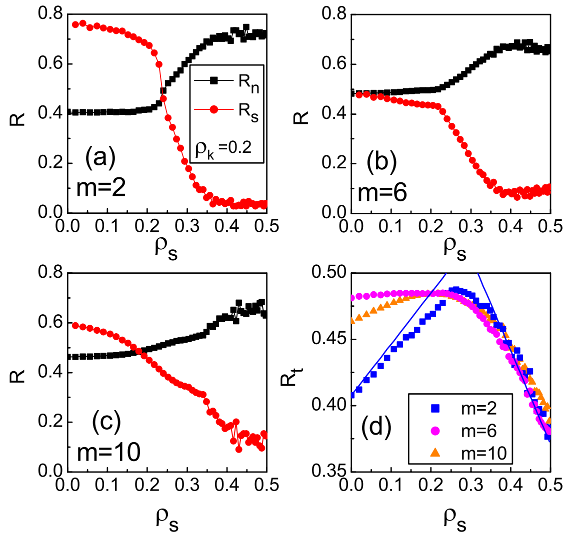

FIG. 3. (Color online) The success rates of the normal agents (squares) and spies (circles) as a function of the fraction of spies ρ_s for (a) $m=2$, (b) $m=6$, and (c) $m=10$. (d) The success rate R_t of the whole population, regardless of the type of agents, as a function of ρ_s for the three values of m.

Therefore when there are only a few spies, they can really exploit the information provided by the normal agents. For small m, the normal agents have a severe crowd effect, leading to a small success rate of about 0.4. This leaves some winning quotas unused. By avoiding the crowd of normal agents, the few spies win more readily. Note that even if all the spies act the same way, the winning quotas left by the normal agents and created by the spies themselves are sufficient to accommodate them when the number of spies is small. However, as the fraction of spies increases to $\rho_s \sim 0.25$, R_n increases and R_s drops abruptly. Too many spies in the population create a crowd effect of the spies. For $m=2$, while the spies know of the minority decision of a k group of normal agents, they may decide collectively to choose a particular decision. The decision would have been the winning decision, without counting the spies. With many spies, however, the spies' decisions may jeopardize their own success rate as their actions turn the would-be winning decision into a losing majority decision. Accompanying with this effect is that the winning history bit string is also changed by the presence of the spies. This actually helps the normal agent, as the history bit string in the basic MG affects the strategies' performance and leads to the severe crowd effect. As a result, R_s drops as ρ_s becomes large. The majority would-be losers in the normal agents become winners and R_n increases. For $\rho_s=0.5$, R_s drops to nearly zero indicating that all the spies are losers and R_n increases to a value above 0.7. For $m=6$ [see Fig. 3(b)], the normal agents in the original MG have a higher success rate of $R_n \simeq 0.46$. This makes the spies harder to win through their collected information, as different k groups will suggest different minority decisions; and there are not too many winning quotas left by the normal agents. More spies will lead to a crowd effect of the spies and actually make room for more normal agents to win. Thus the spies lose the game more often than normal agents on average. For $m=10$ [see Fig. 3(c)], the original MG is in the random coin-toss limit. The information provided to the spies is not useful. For small ρ_s, both the spies and the normal agents make use of the unused winning quotas. For more spies, their presence usually slips what would be the winning (losing) action into the losing (winning) action. Their success rate R_s thus drops with their number. Their presence makes room for more winners and the normal agents take advantage. As a result, R_n increases with ρ_s. Spies will succeed only when there are a few of them. Too many spies create a crowd effect among the spies that benefits the normal agents.

To study the performance of the system as a whole, the success rates R_t of the whole population, regardless of the types of agents, for $m=2$, 6, and 10 are shown in Fig. 3(d). R_t exhibits a peak at some value of ρ_s, which depends on the value of m. This feature can be explained as follows. Since the spies randomly choose k ($k \ll N_n$) normal agents to collect information, their success rate then is related to the success rate R_n of the normal agents. Therefore we may express R_s as a function of R_n of the following form:

$$R_s(m,\rho_s) \equiv F(R_n(m,\rho_s)) = \sum_{i=0}^{k/2-1} C_k^i (R_n(m,\rho_s))^i [1-R_n(m,\rho_s)]^{k-i} + \frac{1}{2} C_k^{k/2} (R_n(m,\rho_s))^{k/2} [1-R_n(m,\rho_s)]^{k/2}, \quad (1)$$

when k is an even number and

$$R_s(m,\rho_s) \equiv F(R_n(m,\rho_s)) = \sum_{i=0}^{(k-1)/2} C_k^i (R_n(m,\rho_s))^i [1-R_n(m,\rho_s)]^{k-i}, \quad (2)$$

when k is an odd number. The first term of the right hand side of Eq. (1) is the probability for a spy selecting the would-be winning decision predicted by the normal agents after he scouts the k group of normal agents. The second term is the probability for a spy selecting the would-be winning decision when the k group shows a tie. The right hand side of Eq. (2) shows the case when k is an odd number. The success rate R_t can be obtained by

$$R_t(m,\rho_s) = \rho_n R_n(m,\rho_s) + \rho_s R_s(m,\rho_s). \quad (3)$$

Substituting of Eqs. (1) or (2) into Eq. (3), we get

$$R_t(m,\rho_s) = \rho_n R_n(m,\rho_s) + \rho_s F(R_n(m,\rho_s)). \quad (4)$$

When ρ_s is small, we assume that $R_n(m,\rho_s) \approx R_n(m,0)$, which is the success rate in the basic MG. The total success rate $R_t(m,\rho_s)$ will show a linear behavior with the increase of ρ_s. Substituting $R_n(m,0)$ into Eq. (4), we get

$$R_t(m,\rho_s) \approx \rho_n R_n(m,0) + \rho_s F(R_n(m,0)), \quad (5)$$

for small ρ_s. As an example, we plot the case of $m=2$ [left solid line in Fig. 3(d)], which fits well to the simulation data by taking $R_n(m=2,0)=0.406$.

When ρ_s is large, $R_n(m,\rho_s)$ takes on another constant value because the many spies often choose the losing action.

Thus Eq. (4) predicts another linear behavior when ρ_s is large. As an example, the right solid line in Fig. 3(d) is plotted by substituting $R_n(m=2,0.5)=0.724$, which is obtained in Fig. 3(a), into Eq. (4). The results agree with the total success rate obtained numerically.

An estimate of the value of ρ_s, where the total success rate peaks, can be found by solving the following equation:

$$\rho_n R_n(m,0) + \rho_s F(R_n(m,0) = \rho_n(1 - R_n(m,0)) + \rho_s(1 - F(R_n(m,0)), \quad (6)$$

which implies that the numbers of agents choosing the two possible decisions are nearly the same. The solution gives

$$\rho_s = -\frac{1 - 2R_n(m,0)}{1 - 2F(R_n(m,0\ )}\rho_n. \quad (7)$$

Substituting $R_n(m=2,0)=0.406$, $R_n(m=6,0)=0.481$, and $R_n(m=2,0)=0.463$, which are obtained from the simulation data, into Eq. (7), we get $\rho_s^*=0.24$, 0.22, and 0.22 for $m=2$, 6, and 10, respectively. Comparing to the simulation results, where $\rho_s^*=0.25$, 0.23, and 0.21 for $m=2$, 6, and 10, respectively, we find Eq. (7) is a good approximation to describe the peak character when ρ_s increases.

IV. SUMMARY

In summary, we studied the minority game in an inhomogeneous population consisting of normal agents and some spies. The spies do not carry any strategy. Instead, they decide their action by scouting a selection of k normal agents and take the minority action of the k group. For a few spies and small values of m, there is useful information by scouting the normal agents and the spies attain a higher success rate than the normal agents. For an intermediate value of m (near $m=6$), the information of the normal agents becomes less useful and the crowd effect of the spies hurts the success rate of the spies themselves. As a result, the normal agents have a higher success rate than the spies. More spies actually assist more normal agents to win, as the spies also provide more winning quotas. When half of the population are spies, the success rate of the spies becomes very low, for the whole range of m studied. The total success rate over the population shows a nonmonotonic behavior as a function of the fraction of spies in the population.

ACKNOWLEDGMENTS

This work was supported in part by the NNSF of China under Grant No. 10635040. P.M.H. acknowledges the support from the Research Grants Council of the Hong Kong SAR Government under Grant No. CUHK-401005. One of us (Y.Y.Y.) would like to thank Professor Da-fang Zheng for helpful discussions.

[1] D. Challet and Y. C. Zhang, Physica A **246**, 407 (1997).
[2] D. Challet, M. Marsili, and Y. C. Zhang, Physica A **294**, 514 (2001).
[3] D. Challet, M. Marsili, and R. Zecchina, Phys. Rev. Lett. **84**, 1824 (2000).
[4] N. F. Johnson, P. Jefferies, and P. M. Hui, *Financial Market Complexity* (Oxford University Press, Oxford, 2003).
[5] See, for example, the articles in *The Complex Networks of Economic Interactions*, edited by A. Namatame, T. Kaizouji, and Y. Aruka (Springer-Verlag, Berlin, 2006).
[6] W. B. Arthur, Pap. Proc. Annu. Meet. Am. Econ. Assoc. **84**, 406 (1994).
[7] R. Savit, R. Manuca, and R. Riolo, Phys. Rev. Lett. **82**, 2203 (1999).
[8] D. F. Zheng and B. H. Wang, Physica A **301**, 560 (2001).
[9] R. Savit, R. Manuca, and R. Riolo Phys. Rev. Lett. **82**, 2203 (1999).
[10] R. Manuca, Y. Li, R. Riolo, and R. Savit, Physica A **282**, 559 (2000).
[11] N. F. Johnson, M. Hart, and P. M. Hui, Physica A **269**, 1 (1999).
[12] M. Hart, P. Jefferies, N. F. Johnson, and P. M. Hui, Physica A **298**, 537 (2001).
[13] T. S. Lo, H. Y. Chan, P. M. Hui, and N. F. Johnson, Phys. Rev. E **70**, 056102 (2004).
[14] T. S. Lo, K. P. Chan, P. M. Hui, and N. F. Johnson, Phys. Rev. E **71**, 050101(R) (2005).
[15] K. F. Yip, P. M. Hui, T. S. Lo, and N. F. Johnson, Physica A **321**, 318 (2003).
[16] M. Anghel, Z. Toroczkai, K. E. Bassler, and G. Korniss, Phys. Rev. Lett. **92**, 058701 (2004).
[17] N. F. Johnson, P. M. Hui, D. F. Zheng, and M. Hart, J. Phys. A **32**, L427 (1999).
[18] K. F. Yip, T. S. Lo, P. M. Hui, and N. F. Johnson, Phys. Rev. E **69**, 046120 (2004).
[19] Y. B. Xie and B. H. Wang, Eur. Phys. J. B **47**, 587 (2005).
[20] L.-X. Zhong, D. F. Zheng, B. Zheng, and P. M. Hui, Phys. Rev. E **72**, 026134 (2005).
[21] S. Moelberta and P. De Los Rios, Physica A **303**, 217 (2002).
[22] E. Burgos and H. Ceva, Physica A **337**, 635 (2004).
[23] T. Kalinowski, Hans-Jorg Schulz, and M. Bries, Physica A **277**, 502 (2000).