Fordham University
Department of Economics
Discussion Paper Series

Endogenous Neighborhood Selection and the Attainment of Cooperation in a Spatial Prisoner's Dilemma Game

Jason Barr
Rutger University, Newark, Department of Economics
Troy Tassier
Fordham University, Department of Economics

Discussion Paper No: 2008-21
October 2008

Department of Economics
Fordham University
441 E Fordham Rd, Dealy Hall
Bronx, NY 10458
(718) 817-4048

Endogenous Neighborhood Selection and the Attainment of Cooperation in a Spatial Prisoner's Dilemma Game*

Jason Barr
Department of Economics
Rutgers University, Newark
jmbarr@rutgers.edu

Troy Tassier†
Department of Economics
E528 Dealy Hall
Fordham University
Bronx, NY 10458
tassier@fordham.edu

September 8, 2008

Abstract

There is a large literature in economics and elsewhere on the emergence and evolution of cooperation in the repeated Prisoner's Dilemma. Recently this literature has expanded to include cooperation in spatial prisoner dilemma games where agents play only with local neighbors in a specified geography. In this paper we explore how the ability of agents to move and choose new locations and new neighbors influences the emergence of cooperation. First, we explore the dynamics of cooperation by investigating agent strategies that yield Markov transition probabilities. We show how different agent strategies yield different Markov chains which generate different asymptotic behaviors in regard to the attainment of cooperation. Second, we investigate how agent movement affects the attainment of cooperation in various spatial networks using agent based simulations.

Key words: repeated prisoner's dilemma, cooperation, agent-based economics, endogenous networks, Markov chains

JEL Classification: C63, C72, C73, D85

*Versions of this paper were presented at the 2008 Eastern Economic Association Meetings, Boston, and at the 14th International Conference on Computing in Economics and Finance, June 2008, Paris. We thank the seminar participants for their helpful comments. We would also like to think Allen Wilhite for his insightful comments. Any errors belong to the authors.

†Corresponding Author

1 Introduction

The repeated Prisoner's Dilemma game, though widely studied, continues to generate interest among economists and social and natural scientists, in general. Despite the pessimistic outcome of the one-shot game, which has only one Nash equilibrium with both agents defecting, the repeated game offers a veritable embarrassment of riches in terms of equilibrium outcomes.

In addition, in both the human and natural worlds, though many opportunities for selfish behavior constantly arise, agents often display cooperative behavior (Poundstone 1992). As the folk theorem indicates, almost any infinite series of actions is a possible equilibrium, given that defect-every period is an available strategy. The folk theorem, however, says nothing about what kind of behaviors may emerge that can foster a long-run cooperative equilibrium. To address the issue of equilibrium emergence several approaches have been taken including evolutionary game theory (Fudenberg and Maskin, 1990), equilibrium selection models (Imhof, et al., 2005) and also agent-based approaches (Wilhite, 2006).

A variation on the theme of the RPD is to incorporate a spatial dimension to agent interactions since players must "live" somewhere, and they generally interact with their neighbors within larger networks of say cities, organizations, or social networks. The spatial Prisoner's Dilemma has received wide attention within theoretical biology (Nowak et al., 1994), but less so in economics. In addition, economic agents can decide where to locate, and this choice may depend on the actions of their neighbors (rivals). Thus it is natural to embed a location choice of agents into a spatial Prisoner's Dilemma.

In this paper we focus on outcomes of a Prisoner's Dilemma game where agents choose their neighborhood and thus choose their interaction partners. We allow agents to compare the payoffs of their current location to other locations, and they can move if the payoffs in the alternative location are greater.

There are competing theories as to the effect of agent movement on the attainment of cooperation in a spatial Prisoner's Dilemma. Agents receive larger payoffs if they live in a neighborhood with large numbers of cooperating agents (this is true whether the agent of interest is a defector or cooperator herself). Thus agents currently located in a neighborhood with many defectors can improve their payoffs by moving to a neighborhood with more cooperating agents. In essence, the ability to leave defectors allows agents an indirect means of punishing defection by opting out of low payoff interactions. Competing with this theory is the idea that agent movement gives defecting agents the ability to invade cooperating neighborhoods. This invasion may deteriorate levels of cooperation in the neighborhood and lead to lower levels of cooperation overall. We focus on these two competing possibilities in this paper. We discuss the conditions that generate the ability of agent movement to increase cooperation and provide simulation examples which demonstrate these conditions. To the best of our knowledge, no other work has explored the roles that neighborhood comparison and movement can have on cooperation.

In summary, we show that the effect of movement is a function of network and neighborhood structure. Here we explore three different networks: a circle, a lattice, and "discrete" networks, which are several separate networks with agents fully-connected within each network. We show, for example, that with a circular network, movement can greatly increase cooperation, while there is a non-monotonic relationship between cooperation and neighborhood size. With lattice networks, we again see a very large increase in cooperation with movement. But the effect is larger when neighborhoods have 8 neighbors, rather than 4 neighbors. However, with the discrete networks, movement, in fact, decreases cooperation. We discuss the reason for these findings below.

The rest of this paper proceeds as follows. The next section gives an overview of the literature on cooperation in the Prisoner's Dilemma, spatial PD games, and agent movement. Then section 3 introduces the basics of the PD model. Following that, section 4 introduces the set of strategies that agents play. Here we focus on a specific class of mixed-strategies that have the Markov property. Section 5 discusses the effects of agent movement within a spatial setting. Then section 6 presents the results of our simulation experiments, which looks at the effect of agent movement in three different network structures. Lastly, section 7 offers some concluding remarks.

2 Related Literature

Studies of the maintenance of cooperation in the Prisoner's Dilemma are vast. Economists are long familiar with the "folk theorem" (Fudenberg and Tirole, 1996), which says that agents can maintain cooperation in an infinitely repeated Prisoner's Dilemma as long as the future is not discounted too heavily.

More recently, other means of maintaining cooperation in the Prisoner's Dilemma have been studied. Examples range to include reputation (Nowak and Sigmund, 1998), reciprocity (Axelrod, 1984), the use of tags and signals to recognize opponent types (Riolo, 1997; Riolo et al., 2001; Hales, 2001; Janssen, 2008), and withdrawal from play (Aktipis, 2004; Janssen, 2008).

Previous research has shown that incorporating a spatial element into a repeated Prisoner's Dilemma setting can result in cooperation due to the assortment of neighborhoods into areas of cooperation and areas of defection (Brauchli et al. 1999, Ferriere and Michod 1995, Killingbach and Doebeli 1996, Nowak and May 1992). This research suggests that spatial interactions are one way to produce repeated interactions that can result in cooperative behavior. In this paper we extend this literature by having agents move or relocate in the space in which interactions take place; agent strategies are governed by behavioral rules commonly used in spatial studies of the Prisoner's Dilemma.

Again note that the results derived from movement are not obvious at the outset. Movement may allow agents with a propensity to defect the ability to invade cooperators and thus lead to lower levels of cooperation. Or movement may allow agents with a propensity to cooperate the ability to avoid defectors and lead to higher levels of cooperation. Thus the focus of this paper is on how movement affects the attainment of cooperative outcomes (equilibria) within the behavioral structure described below. Finally, the ability of agents to move across a spatially defined interaction structure may be seen as an example of endogenous network formation in repeated play games.

Most closely related to the implementation of the Prisoner's Dilemma in our paper is the work on the maintenance of cooperation in spatial models (Nowak and May, 1992; Nowak, Bonhoeffer and May, 1994; Schweitzer, Behera and Muhlenbein, 2002). In these models it is shown that repeated interaction with local neighbors may lead to the evolution of cooperation in a Prisoner's Dilemma.[1] Wilhite (2006) explores RPD cooperation for several different network architectures. Agents update their actions each period by imitating their rivals, playing the same action as the rival with the largest payoff in the previous period. He finds that in a complete network (with synchronous updating), all agents become defectors after the first round of play, as long as one agent was a defector initially. On the other hand, with a star network, all agents play the same action after one round, but it can be either cooperate or defect; which strategy emerges in equilibrium is a random variable. With a ring, on the other hand, both actions can exist side by side in equilibrium, but cooperators tend to make up the majority of actions. With a grid (lattice) again, cooperators and defectors can live side by side, in equilibrium, but the dynamics can often behave chaotically. In sum, Wilhite's paper shows with even a simple imitation-based rule, network structure can have important implications for the emergence of cooperation, and we explore this effect with both agent movement and mixed-strategies.

There has been a limited amount of work on movement in the spatial Prisoner's Dilemma. Aktipis (2004) studies the performance of a behavioral rule named "Walk Away" (WA). With WA an agent always cooperates but if its opponent defects, the agent moves to a new

[1] The introduction to Schweitzer, Behera and Mühlenbein (2002) provides a nice overview of the literature on cooperation in the Prisoner's Dilemma.

location. This strategy is tested in a tournament environment similar to Axelrod (1984). The WA strategy is tested against standard rules such as Anatol Rapport's Tit-for-Tat, (copy opponent's last action) and Nowak and Sigmund's PAVLOV (switch action from cooperate to defect or defect to cooperate if the agent's opponent defects.) Aktipis shows that the simple walk away strategy performs well against these other well-known strategies even though it is very simple. WA is successful because movement allows agents to engage in repeated interactions with cooperators but avoid defectors.

Note however that movement may not necessarily lead to cooperation in general. Movement could be valuable for agents with a propensity to defect by allowing them the ability to invade and exploit a neighborhood of cooperators (Dugatin, 1992; Dugatin and Wilson, 1991; Enquist and Leimar, 1993). Thus the study of movement within the context of standard economic behavioral rules and across different spatial structures is warranted.

The literature on exit or refusal to play may be seen in a similar spirit to movement (Janssen, 2008; Schluessler, 1989; Vanberg and Congelton, 1992). In our context agents who refuse to play with a past defector may be seen as agents who no longer play because they have re-located.

In a similar spirit to our work, Hanaki et al. (2007) investigate the evolution of cooperation when agents can decide whether to create or destroy network connections based on the costs and benefits of these connections (i.e., they present a model of partner choice). Again, as is common in the literature, agents imitate the action of the rival with the highest payoff. They find, for example, that there is a positive correlation between the size of the network and the degree of cooperation. In addition, they also find that "sparse" networks are needed to maintain cooperation; that is, agents cannot be fully connected for cooperation to be sustained.

		Rival's Action	
		Coop.	Defect
Agent's Action	Coop.	A	B
	Defect	$C = A + \varepsilon$	$D = B + \mu$

Table 1: An agent's payoffs in the Prisoner's Dilemma game

3 The PD Game on a Network

In this paper, we work with the standard Prisoner's Dilemma game, whose payoffs structure is given in Table 1. where $C > A > D > B$ and $C + B < 2A$. To simplify matters we set the "cheat" payoff such that $C = A + \varepsilon$, where $\varepsilon > 0$. Define an agent and his rival's actions as $x, y \in \{0, 1\}$, where $x, y = 1$ if an agent (rival) cooperates, 0 otherwise; let $\boldsymbol{\alpha} \equiv \{A, B, C, D\} = \{A, B, A + \varepsilon, B + \mu\}$. Then we can write the payoff function for any agent, given a rival's action as

$$\begin{aligned}\pi(x, y; \boldsymbol{\alpha}) &= Axy + Bx(1 - y) + C(1 - x)y + D(1 - x)(1 - y) \\ &= Ay + B(1 - y) + [\varepsilon y + \mu(1 - y)](1 - x).\end{aligned}$$

Further, suppose an agent plays against n rivals, the average payoff is given by

$$\bar{\pi}(x, y; \boldsymbol{\alpha}) = \frac{1}{n_x}\frac{1}{n_y}\sum_{i=1}^{n} \pi(x_i, y_i) = A\rho + B(1 - \rho) + [\varepsilon\rho + \mu(1 - \rho)](1 - \lambda) \equiv \pi(\lambda, \rho), \quad (1)$$

where $\rho = \sum y_i / n_y$ is the proportion of rivals that play cooperate, and $\lambda = \sum x_i / n_x$ is the proportion of times an agent cooperates against her neighbors. In short, the average payoff to the agent is given by a weighted sum of A and B payoffs, plus a "defector's bonus" given by $[\varepsilon\rho + \mu(1 - \rho)]$, which is earned in $100(1 - \lambda)\%$ of games against the rivals. Notice that $\partial\pi(\lambda, \rho)/\partial\lambda = -[\varepsilon\rho + \mu(1 - \rho)] < 0$, and thus, *ceterus paribus*, an increase in the cooperation rate for an agent will reduce her payoff. If everyone cooperates then the agent gets an average payoff of A; if everyone defects the agent gets an average payoff of $B + \mu = D$.

Also notice that $\partial\pi(\lambda,\rho)/\partial\rho = (A-B)+(\varepsilon-\mu)(1-\lambda) > 0$, which means an agent will strictly prefer to play with more cooperators, if given the choice of with whom to play.[2]

4 Markov Strategies

We focuses on a class of mixed strategies that evolve over time based on an agent's play with his rivals (similar in spirit to Nowak, et al., 1994). Here each agent's probability of playing cooperate against a rival in the next round is given by

$$p_{t+1} = f(\lambda_t, \rho_t; \boldsymbol{\alpha}),$$

where λ_t is the proportion of times an agent cooperates against his rivals, and ρ_t is the fraction of rivals who cooperate in a particular round; $1-p_{t+1}$ is the agent's probability of defecting. Denote $f(\lambda_t, \rho_t; \boldsymbol{\alpha})$ as the *Probability Evolution Function* (PEF). Note that here a "round" is the time during which all agents play against each of their rivals. For example, if there are N agents and each agent has n rivals, then a round constitutes a total of Nn games, with a total of $2Nn$ action choices (assuming each agent chooses n actions according to his PEF; and n rivals each choose an action according to their PEFs).[3]

In short, the strategy for an agent is evolving based on the outcomes of last round's play with his rivals. For the moment, we make no assumptions about the PEF, except simply that $p_t \in [0,1]$ for all rounds and all agents. Further we assume that the PEF is determined by play with agents that are directly in an agent's network; for example, if an agent's network is described as a particular graph, the PEF is based on play with rival's to whom the agent is directly connected in the graph.

[2] This assumes that if $\varepsilon < \mu$, then $(A-B) > (\varepsilon - \mu)$.

[3] For the rest of this section, for simplicity, we assume that each agent has the same number of rivals. In the simulations given below, the number of rivals may not be the same for all agents, depending on the network structure. Also, when clarity can be preserved, time subscripts will be dropped for notational convenience.

This class of strategies generates a Markov chain, where the state of the system is described by a value k. Depending on the specific PEF, k can either be the number of cooperators in a given round, or it can be an integer that indexes the (uniquely determined) outcomes from a round of play. Denote the set $\mathbf{q}_{it} \equiv \{x_{i1t}, ..., x_{int}, y_{i1t}, ..., y_{int}\} \in \{0,1\}^{2n}$ as the set of actions by agent i and his rivals in a particular round t. Then we can denote $\mathbf{Q}_t = \times_{i=1}^{N} \mathbf{q}_{it}$ as the set of all actions played in a given round. Each set $\mathbf{q}_{it}$ has $(2n)^2$ possible outcomes, and $\mathbf{Q}_t$ has $(2Nn)^2$ possible outcomes. Thus we have the index set $k \in \left\{0, 1, 2, ..., (2Nn)^2 - 1\right\}$, which assigns an integer to each possible outcome in $\mathbf{Q}_t \in \mathbf{Q}$, where $\mathbf{Q}$ is the set of all binary vectors of size $2Nn$.

It is straightforward to show that the PEF generates a Markov chain. Each round, we denote a realization of actions as $\mathbf{Q}_t \in \mathbf{Q}$; this vector of actions then gives rise to each agent's updated PEF. That is, $p_{it+1} = f\left(g\left(\mathbf{q}_{it}\right)\right)$, were $g(\cdot)$ converts actions to their respective cooperation rates. In addition, k can be generated from $\mathbf{Q}$ (such as when k is the base-10 integer equivalent of $\mathbf{Q}_t$), i.e., $k_t = \sum_{i=1}^{2Nn} z_{it} 2^{2Nn-i}$, where $z_{it} \in \{0,1\}$ is the i^{th} element of $\mathbf{Q}_t$.

Given that we have a state k_t, and each agent has a probability of playing cooperate in the next round, p_{it+1}, we can (with some effort) determine the probability of each outcome of $\mathbf{Q}_{t+1}$ and hence k_{t+1}. Thus the probability of going from state k_t to k_{t+1} is a function of $\mathbf{Q}_t$, which is determined by the vector of probabilities $\{p_{1t}, ..., p_{Nt}\}$.

In summary, each agent has a probability of cooperating. This gives rise to a distribution for k, the number of possible states. Given that each k is associated with a particular outcome of a round of the RPD, each Q_t and hence k_t gives rise to a set of probabilities for each agent's PEF. Thus each k let's us calculate the probability of going to state $k+1$. This is now stated formally.

Proposition 1 *A RPD game, where each agent plays a mixed strategy according to* $f\left(\lambda_t, \rho_t; \boldsymbol{\alpha}\right)$, *gives rise to a Markov chain, with transition probabilities given by*

$$
\begin{aligned}
T(k_{t+1}|k_t) &= \Pr(K = k_{t+1}|k_t) = h(f(\lambda_{1t}, \rho_{1t}), ..., f(\lambda_{Nt}, \rho_{Nt})) \\
&= h(p_1(k), ..., p_N(k)),
\end{aligned}
$$

where $h(\cdot)$ is a function which converts the vector of agents' cooperation probabilities to a distribution over k.

Proof. Each round gives rise to a realization of actions, $\mathbf{Q}_t = \times_{i=1}^{N}\mathbf{q}_{it}$. This realization can then be assigned an index number, $k_t = \sum_{i=1}^{2Nn} z_{it}2^{2Nn-i}$, where z_{it} is the i^{th} element of $\mathbf{Q}_t$. Given $\mathbf{Q}_t$, each agent updates his probability of playing cooperate based on $p_{it+1} = g(\mathbf{q}_i) = f(\lambda_i, \rho_i)$. p_{it+1} then gives rise to a distribution over k_{t+1}, since $Prob(\mathbf{q}_{it+1} = \mathbf{q})$ is a function of probabilities that each agent (and rival) will cooperate in any given game. Thus, the RPD game gives rise to Markov transition probabilities: $T(k_{t+1}|k_t) = \Pr(K = k_{t+1}|k_t)$, since the probability that the state will be k_{t+1} is determined by k_t, which emerges from the realization of $\mathbf{Q}_t$. ■

Note that the function $h(\cdot)$ can be simple or complicated based on the spatial topology on which the agents interact. Once the transition matrix is given and initial probabilities of playing cooperate are assigned, the asymptotic probabilities of being in a given state, $\tau^t = T\tau^0$, can be determined, where τ^t is the probability of being in each state at time t. In other words, we can investigate the outcome of the repeated Prisoner's Dilemma game based on the properties of the Markov chain generated by the probability evolution function.

4.1 Example I

Suppose that we have $N = 3$ agents, who are each connected to each other. Further suppose that each agent determines his probability of playing cooperate simply based on what his other two rivals do, and which is given by $p_i = \frac{A\kappa_i}{2A+\varepsilon(2-\kappa_i)}$, where $\kappa_i \in \{0, 1, 2\}$ is the number of rivals that cooperate. To make things simple, assume that each agent chooses one action at the beginning of the round and plays the same action against all rivals. Notice that this one action assumption reduces the number of possible states.

For example, let a set of actions be given by $\mathbf{x}_t = \{1,1,0\}$, then $\mathbf{p}_{t+1} = \left\{\frac{A}{2A+\varepsilon}, \frac{A}{2A+\varepsilon}, 1\right\}$. The transition matrix is then given by $\mathbf{T} =$

		k_{t+1}: 3	2	1	0
k_t	3	1	0	0	0
	2	$\left(\frac{A}{2A+\varepsilon}\right)^2$	$2\left(\frac{A}{2A+\varepsilon}\right)\left(\frac{A+\varepsilon}{2A+\varepsilon}\right)$	$\left(\frac{A+\varepsilon}{2A+\varepsilon}\right)^2$	0
	1	0	$\left(\frac{A}{2A+\varepsilon}\right)^2$	$2\left(\frac{A}{2A+\varepsilon}\right)\left(\frac{A+\varepsilon}{2A+\varepsilon}\right)$	$\left(\frac{A+\varepsilon}{2A+\varepsilon}\right)^2$
	0	0	0	0	1

In this case, each row of $\mathbf{T}$ is binomial distribution. Notice that there are two absorbing states, meaning that in this case, all agents defecting or all agents cooperating are the only asymptotic outcomes.

4.2 Example II

Assume now that there are only two agents, and further each agent's probability is a weighted average of his action and the action of his rival, i.e., $p_i = \beta x_i + (1-\beta) x_{-i}$, where $\beta \in (0,1)$ Here the action set $\mathbf{Q} = \{\{0,1\}, \{0,1\}\}$, with 4 possible states. Thus the transition matrix is $T =$

$$\begin{array}{ccc|cccc|}
 & & & \multicolumn{4}{c}{k_{t+1}} \\
 & & & \{1,1\} & \{1,0\} & \{0,1\} & \{0,0\} \\
 & & & 3 & 2 & 1 & 0 \\
\hline
 & \{1,1\} & 3 & 1 & 0 & 0 & 0 \\
k_t & \{1,0\} & 2 & \beta(1-\beta) & \beta^2 & (1-\beta)^2 & \beta(1-\beta) \\
 & \{0,1\} & 1 & \beta(1-\beta) & (1-\beta)^2 & \beta^2 & \beta(1-\beta) \\
 & \{0,0\} & 0 & 0 & 0 & 0 & 1 \\
\hline
\end{array}.$$

Notice that in this case, the number of states can be reduced to simply the total number of times cooperation is played in a given round. Then the transition matrix becomes $T =$

$$\begin{array}{cc|ccc|}
 & & \multicolumn{3}{c}{k_{t+1}} \\
 & & 2 & 1 & 0 \\
\hline
 & 2 & 1 & 0 & 0 \\
k_t & 1 & 2\beta(1-\beta) & (1-\beta)^2+\beta^2 & 2\beta(1-\beta) \\
 & 0 & 0 & 0 & 1 \\
\hline
\end{array}.$$

Note that if this rule was applied to three agents, where each agent chooses one action each round and say $p_i = \beta_i x_i + \beta_j x_j + \beta_k x_k$, where $\sum \beta_i = 1$, then the action set is $\mathbf{Q} = \{0,1\}^3$, with 8 possible states, which then could be reduced then be reduced to 4 states, which represents the number of agents who play cooperate each round. Then with N agents, $\mathbf{Q}$ would have N^2 possible actions (states), which could be reduced to $(N+1)$ states. Thus as we expand the number of agents, even for a simple rule the number of states and hence the transition matrix soon becomes quite cumbersome to work with. Note that with a less than fully connected network and with agents playing a new action against each rival, calculating the transition matrices will be even more cumbersome since each agent's probability of cooperating will depend on its neighbors' actions. Thus generating the probabilities of going from one state to another can be quite tedious; for this reason we rely on simulations to explore the likelihood of cooperation emerging.

5 Agent Movement and Spatial Interactions

We now consider the effects of agent movement given their spatial interactions. Throughout the discussion, we assume that each agent is given an opportunity to move to a new randomly chosen location (specified below) with an exogenously given probability. When the opportunity arrives, agents choose to move to the alternative location or to remain in their current location using a myopic decision rule. Specifically, the agent compares the average payoff received in the current round across games she played at the current location to the average payoffs she would have received if she was at the alternative location. To determine these alternative payoffs the agent plays a fictional game with each potential new neighbor according to the neighbors' and agent's current mixed strategies. The agent then calculates an average payoff in these fictional games. If the alternative location would have provided better average payoffs, she moves. Otherwise she stays at her current location.

For a particular agent, movement will occur when $\Delta\pi \equiv \pi(\lambda', \rho') - \pi(\lambda, \rho) > 0$, where λ', ρ' are the cooperation rates in the new, potential location. Since $E[\lambda'] = E[\lambda] = p$, it's straightforward to show that

$$E[\Delta\pi] = E(\rho' - \rho)[(A - B) + (\varepsilon - \mu)(1 - p)].$$

Thus, on average, movement will occur when $(\rho' > \rho)$; that is when the new location has more cooperators than the old.[4]

[4] Note that this is what is expected on average, but it is still possible for an agent to move to the alternative neighborhood when there are fewer cooperators. For this to occur, the agent must randomly choose to defect a sufficiently larger number of times when playing the fictional game in the new neighborhood compared to the number of times she defected in the current neighborhood, so that the extra payoff from defection "pays" for the "loss" of having fewer neighbors who cooperate. In other words, if the agent finds herself, by chance, to be an extreme defector in the alternate neighborhood, she may expect it to be profitable to move into this neighborhood, which could potentially cause defection to spread throughout the network.

5.1 The Effect of Movement

We now want to consider what effect agent movement will have on the attainment of cooperation given the PEF of the agent. We know from above that, on average, an agent will move into the alternative neighborhood if that neighborhood contains more cooperating agents. In order for agent movement to have a positive effect on the attainment of cooperation, movement must result in a transition matrix more favorable for cooperation. However this is a complex process; when an agent moves, the agent effects his own probability of cooperation in the next round as well as the probability of cooperation for the agents in his old (by his exit) and new (by his arrival) neighborhoods. Each of these probabilities directly depends on the specific PEF of the agents and the structure of agent interactions.

To simplify matters we focus only on one specific PEF (equation (2) below); this way we can isolate the effects of movement and network structures. We have specifically chosen a PEF that has two properties that aid in the analysis. The first property is that the PEF has only two absorbing states (equilibria): everyone on the network will either cooperate or defect. Each agent reaches a probability where she always cooperates or always defects in equilibrium, there are no mixed probabilities in equilibrium.[5] This allows us to focus simply on the percent of agents in each absorbing state at the end of a run of a simulation. The second property is that the PEF is imitation based. That is, agents look to the performance of their neighbors when updating their strategies. Imitation is an important part of the evolution of cooperation (Wilhite, 2006; Hanaki, et al., 2007). If all agents simply played a best-response, for example, defection would be the only stable equilibria.

Specifically, we assume that an agent's PEF is given by the relative performance of the

[5] See discussion below.

rivals against him:[6]

$$p = \frac{\sum_j^n \pi\left(y_j, x_j\right) y_j,}{\sum_j^n \pi\left(y_j, x_j\right)}, \tag{2}$$

where x_j is the action choice of an agent against, rival j; y_j is the rival's action choice. Inserting the payoff function (eq. 1) and performing some algebraic manipulation we get the following PEF.

$$p\left(\lambda, \rho\right) = \lambda \frac{A\rho}{A + \varepsilon\left(1-\rho\right)} + \left(1-\lambda\right) \frac{B\rho}{B + \mu\left(1-\rho\right)}.$$

Notice a few properties of this rule: $p\left(\lambda, 0\right) = 0$ and $p\left(\lambda, 1\right) = 1$, which means that if all of the rivals of an agent select the same action, an agent will play that action with probability one. Further the transition matrix will always have everyone cooperate and everyone defect as an absorbing state of the system. Also notice that $\partial p/\partial\rho \geq 0$. This positive partial derivative is an important feature of the PEF. Since, as discussed above, an agent will likely move into a new neighborhood if the percentage of cooperating agents is greater than in her old neighborhood, then the agents probability of cooperation will increase. One may initially think that movement will generate more cooperation with this PEF. But, also consider that an agent with a high likelihood of defection can move into a neighborhood that contains mostly cooperators and decrease all of his neighbor's probability of cooperation. This effect can be more severe when neighborhoods are smaller. And, when this occurs the defecting agent pulls down the probability of cooperation for all of the agents in this neighborhood. This effect can then cascade as each neighbor is then playing with a group of agents who each have a decreased probability of cooperation which can then lead to further decreases in the probability to cooperate, etc. Thus, with this PEF we have the possibility that movement may increase or decrease cooperation.

[6] Our rule is similar in spirit to the one used in Nowak, et al. (1994). However, the rule used here is different in that agents update their probabilities based on their neighbors actions, whereas Nowak, et al.'s rule is based on the actions of neighbors and neighbors' neighbors.

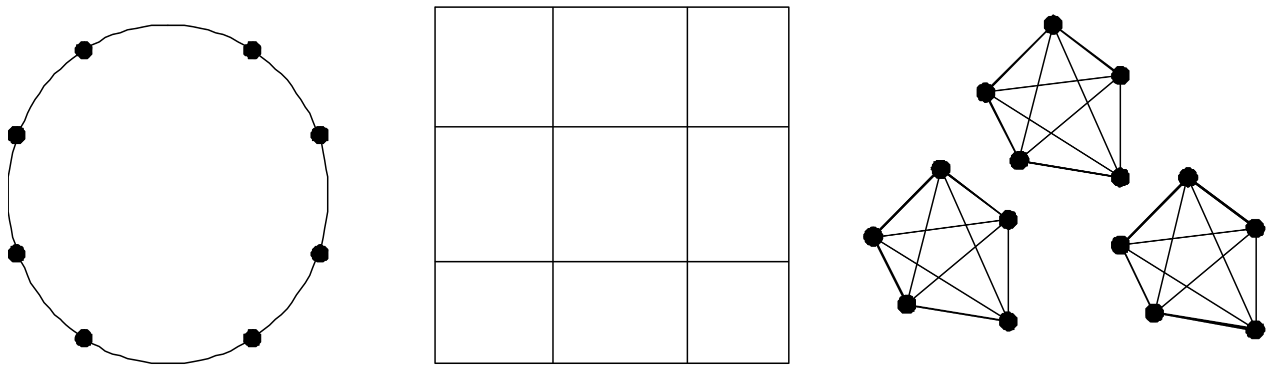

Figure 1: Three types of agent networks: circle, lattice and discrete neighborhoods.

6 Simulations

We now examine how movement and partner selection impacts the attainment of cooperation across various spatial structures. We consider three types of spatial topologies: circular neighborhoods of various size, a lattice based neighborhood, and discrete neighborhoods that are fully connected within a neighborhood but disjoint across neighborhoods. These three are chosen because they allow for a comparison of three common network types, and may be relevant for geographic analysis. Clearly more types of networks can be explored, such as small world or scale-free distributions, but for simplicity we just focus on these three.

In all of the experiments below we will have 100 agents and 144 locations available for agents to locate. Thus 44 locations will be empty in each round and available for agents to move there. We will examine two different scenarios in each parameter set considered: no movement and full movement. With no movement, each agent is assigned to a randomly chosen location at the beginning of the simulation and the agent remains there throughout. With full movement, each round, each agent is allowed the option to move to a randomly chosen vacant location. As described above, if the payoffs the agent received in the current round are less than the payoffs she would have received in the alternative location, she moves to the alternative location; otherwise, she remains in her current location.

Each simulation run proceeds as follows:

1. Initially, each agent is assigned a unique randomly chosen location on a graph (which varies across experiments) and a probability of cooperation randomly drawn from uniform [0,1] distribution.

2. Each agent plays a game with each of his neighbors and the average payoffs are calculated for each agent. Note that for each game, the agent chooses an action according to his probability distribution.

3. When all games have been played for a given round, each agent is then given the opportunity to move to a new location.

4. Once all movement decisions have been made, each agent then updates her p_i according to the results of the various games. Thus the agents update their probability of cooperation simultaneously at the end of a round; mid-round adjustments of the probability to cooperate are not allowed.[7]

5. Then the next round of the simulation occurs, following the same procedure discussed above.[8]

We run the system for 100,000 rounds or until the system reaches an absorbing state (where all agents either cooperate or defect 100% of the time) whichever comes first.[9] These 100,000 rounds are a *run* of the simulation. When a run is complete we calculate the average probability of cooperation across all agents in the population. We complete 100 runs for each

[7] Note that when agents move there is the possibility that there will be new agents, who were not present when the agent played his set of games. The PEF, however, evolves according to who the agent played with and not the new agents, if there are any. Thus in some sense agents' PEF evolve according to the nature of the neighborhood rather than, possibly, the neighbors themselves.

[8] We also have tried several alternative specifications of our model such as: synchronous vs. asynchronous updating of probabilities, different orderings for when an agent updates her probabilities and when she moves, and different initial distributions for the probability of cooperation. None of these alternatives resulted in qualitatively different results.

[9] It is a rare instance that the simulations time-out at 100,000 rounds. But in order to complete the simulations in a timely manner we set this upper bound on the time to converge to an absorbing state.

parameter set and report the average probability of cooperation that results at the end of each run.

6.1 Circular Neighborhoods

We begin with a circular topology where agents are located on a "ring." In this model, a neighborhood is defined by a number of locations f in each direction on the ring; each location has $2f$ neighbors. Enumerate each location in the ring 1, 2,...,144. If an agent is located at location 20 and $f = 1$ then the agent has two neighbors, the agent at location 19 and the agent at location 21 (assuming these locations are not empty.) Note that empty locations may create fewer neighbors for a given agent.

To begin we use a set of base payoff values and vary the number of neighbors f. Payoffs are as follows: $A = 3.0$, $B = 1.0$, $C = 4.0$, and $D = 2.0$. We vary f between 1 and 16 and report the levels of cooperation that result for the no movement and full movement case in Table 2.

Number of Neighbors (2f)	No Movement	Full Movement
2	0.39	1.00
4	0.22	0.46
8	0.07	0.37
16	0.18	0.84
32	0.53	1.00

Table 2: Effect of neighborhood size on probability of cooperation. Percentage of cooperators versus neighborhood size, for a circular network. Results are averages of 100 runs. $A = 3, B = 1, C = 4, D = 2$.

As reported in the table, allowing movement always results in a higher probability of cooperation. The results occur because agents are able to opt out of neighborhoods with agents who commonly defect and move into cooperating neighborhoods. When an agent moves into a neighborhood of cooperators, his probability of cooperation in the next round goes up, this can spark a process of "cooperative tipping." That is to say, because movement raises the probability that agent will cooperate (greater than the negative effect that the agent will have on his neighbors' PEF), the transition probabilities will change to favor moving to a state with more cooperators.

Interestingly, there is a non-monotonic increase in the likelihood of cooperation as the network size increases. For a small neighborhood size it is relatively easy for the agents to coordinate on the all cooperate strategy (see the discussion of this effect in the discrete neighborhood model below.) As the network size increases cooperation becomes more difficult in terms of number of agents who must coordinate on cooperation. But there is a countervailing effect with regard to the network size. With a large network, the agents are more closely aligned/connected across the entire population; a small pocket of cooperators is able to expand their influence further throughout the network. Thus as the network size grows cooperation becomes easier. In fact, when neighborhood size is 64 (32 neighbors on each side), movement generates 100% cooperation.

Next, as a robustness check, we consider the movement effect over a range of payoff values for C leaving the other payoff values unchanged. As C increases it becomes more beneficial for an agent to defect and thus cooperation is more difficult to obtain. We report results for 8 and 16 neighbors in Tables 3 and 4. Again we see that allowing agents to move increases the likelihood of attaining cooperation among agents as long as the payoffs are not too large; if they are too large then no cooperation is possible, with or without movement.

C - Cheat Payoff	No Movement	Full Movement
3.50	0.59	0.97
3.75	0.34	0.79
4.00	0.07	0.37
4.25	0.02	0.00

Table 3: Effect of cheat payoff "C" on probability of cooperation - 8 neighbors.

Overall, the results of this subsection indicate that movement has the effect of substantially increasing the likelihood of cooperative outcomes developing. Agents are able to avoid

C - Cheat Payoff	No Movement	Full Movement
3.50	0.97	1.00
3.75	0.69	0.99
4.00	0.18	0.84
4.25	0.02	0.19

Table 4: Effect of cheat payoff "C" on probability of cooperation - 16 neighbors.

C - Cheat Payoff	No Movement	Full Movement
3.5	0.00	1.00
4.0	0.00	0.97
4.5	0.00	0.58
5.0	0.00	0.03

Table 5: Effect of cheat payoff "C" on probability of cooperation - lattice model, 8 surrounding neighbors.

defectors by opting to move to neighborhoods with higher levels of cooperation. Neighborhood size has a nonmonotonic relationship with cooperations.

6.2 Lattice Based Neighborhoods

Next, we consider an agent network in the form of a 12×12 lattice. Tables 5 and 6 present the results. The first table shows outcomes when each agent's neighborhood is a "Moore" neighborhood, i.e., each agent has 8 surrounding neighbors.[10] The second table presents the results when the neighborhood consists of the 4 surrounding agents (i.e., up, down, left and right). Again, we vary the cheat payoff and investigate the effect of agent movement on the ability of agents to reach cooperative outcomes. As we can see from the tables, especially when agents live within the 8-agent neighborhood, movement has a substantial impact on the evolution of cooperation. However, neighborhood size makes a very strong contribution; with the 4-agent neighborhoods, we can see that the range of cheat payoffs for which movement matters is much smaller than in the Moore neighborhood case.

[10]The edges are not wrapped, so agents on the corners have fewer neighbors.

C - Cheat Payoff	No Movement	Full Movement
3.01	0.00	0.12
3.05	0.00	0.11
3.10	0.00	0.10
3.20	0.00	0.04

Table 6: Effect of cheat payoff “C” on probability of cooperation - lattice model, with 4 surrounding neighbors.

As with the circle network, smaller neighborhoods diminish cooperation. Smaller neighborhood sizes have two effects. One is that each agent’s PEF is, in some sense, more sensitive to living with a defector, since there are fewer of them. Without movement, the asymptotic probabilities from the Markov chain will always have a greater probability for defection; and this probability will increase with smaller neighborhoods. In addition, with smaller neighborhood sizes, there is less of an opportunity for pockets of cooperation to spread.

6.3 Discrete Neighborhood Model

Finally we consider a discrete neighborhood model. In this model, there are M distinct neighborhoods, each with m locations. Each location in a given neighborhood is connected to all other locations in the neighborhood, but there are no connections between neighborhoods. This is a key feature of this model because if cooperation is able to develop in a given neighborhood there is no mechanism that allows cooperation to spill over into the other neighborhoods. This is distinct from the other models discussed above where a small group of cooperating agents can expand outward into other neighborhoods which connect to them either directly or indirectly through neighbors of neighbors.

In addition the effects of agents changing neighborhoods are very direct. An agent moving into a neighborhood takes on a probability of cooperation that matches the other agents in the neighborhood and also directly changes the probability of cooperation for all of the agents in the neighborhood, but there are no indirect effects to other agents elsewhere in the population. In the other two models an agent moving into a neighborhood effects a set of agents but this set is only partially connected to each other. For example if each agent on a ring has two neighbors, one in each direction, the two neighbors only have one of their two neighbors in common. In the discrete neighborhood model all neighbors have the same set of neighbors in common (excepting that the agent is not a neighbor of himself) and no other neighbors. Thus the discrete neighborhood model allows us to parse out some of the effects of the other models listed above.

In these simulations we have 100 agents and 144 total locations ($Mm = 144$). We vary the number of neighborhoods and size over different simulations. We find that cooperation is more difficult to attain in this model than in the other two. Thus we selected payoffs such that cooperation is easier to be attained. Specifically, we again choose $A = 3.0$ and $B = 1.0$ but set C and D to lower values ($C = 3.1$ or $C = 3.05$ and $D = 1.1$ or $D = 1.05$ depending on the simulation run.) Thus the defector's bonus is smaller than in the models discussed above. Results for the discrete neighborhood model are shown in Tables 7 - 9.

There are two key results to take from these tables. The first is that movement decreases the likelihood of attaining cooperation. There are two reasons why this occurs. If cooperation results in a given neighborhood there is not mechanism which allows cooperation to spread since the neighborhoods are disjoint from each other. Thus pockets of cooperation that occur by chance cannot propagate into other regions of the population. In addition, recall that the partial derivative of the PEF with respect to rivals cooperating is positive. Thus when a defecting agent moves into a new neighborhood, her probability of cooperation increases (recall that agents only improve payoffs by moving into neighborhoods with more cooperators). But since she is moving from a neighborhood with fewer cooperators her effect on everyone else in the new neighborhood is negative. In addition, each new neighbor will have a lower probability of cooperation and will be playing with the other neighbors who now each have a lower probability of cooperation. Thus each neighbor will be even less

C	D	No Movement	Full Movement
3.05	1.05	0.07	0.00
3.10	1.10	0.02	0.00

Table 7: Discrete Neighborhood Model - 8 Neighborhoods of size 18.

C	D	No Movement	Full Movement
3.05	1.05	0.23	0.07
3.10	1.10	0.09	0.00

Table 8: Discrete Neighborhood Model - 12 Neighborhoods of size 12.

likely to cooperate in the future and a downward spiral can be created toward an all defect outcome. In the other two models this spiral is less likely to occur because each neighbor has other neighbors not affected by the new agent who anchor them nearer their current probability of cooperation. These other neighbors mitigate the effects of a defector invading and pulling down everyone's probability of cooperation too far. Thus cooperation becomes easier to obtain in the first two models.

A second key result is that it is easier to obtain cooperation when there are a large number of neighborhoods and thus each neighborhood is small. This harkens back to the discussion of the circle model where the smallest neighborhood sizes had high levels of cooperation. Then, as the neighborhood size increased, cooperation first decreased and then increased again. The eventual increase ties in with the mitigating effects of the additional neighbors that were discussed in the previous paragraph. Having more agents helps to anchor regions of cooperating agents to their strategies.

Finally, it is worth re-emphasizing that the PEF used in this paper only produces positive movement effects in some network interaction structures. Thus the ability of movement to produce cooperative outcomes is a function of both the behavioral rules that agents use in playing the game as well as the interaction structure of the agents.

C	D	No Movement	Full Movement
3.05	1.05	0.31	0.21
3.10	1.10	0.20	0.07

Table 9: Discrete Neighborhood Model - 18 Neighborhoods of size 8.

6.4 General Discussion

As demonstrated in this paper, the ability of movement to aid cooperation is a function of both the specific PEF as well as the network structure on which agents locate, and therefore finding general sufficient or necessary conditions for the effect of movement on the attainment of cooperation is probably not possible. However, below we sketch our ideas for some key aspects of the ability of agents to improve the likelihood of cooperation with movement.

Here we demonstrate that if agents can move to open spaces on a network and can choose their location based on a comparison of payoffs, then the nature of the PEF and the agent network can determine whether agent movement increases or decreases propensities to cooperate.[11] Specifically, we offer that if the following properties hold, the ability of agents to move will increase propensities to cooperate:

1. $\partial\pi/\partial\rho > 0$. That is, the higher the cooperation rate of neighbors, the larger the profits for an agent. This is true by definition for the Prisoner's Dilemma. Agents must prefer living with cooperators in order for movement to increase cooperation.

2. Agents can compare payoffs from play with current rivals versus play with a new set of rivals. Agents must be able to make an intelligent decision as to where to locate if given the opportunity to move. Thus there is a minimum bound on the intelligence of agents needed for movement to be beneficial.

3. $\partial p_{t+1}/\partial\rho > 0$. Increasing cooperation by rivals will induce a higher probability of cooperation the next round. This implies that holding constant the number of times

[11] We assume that there are no moving costs.

an agent cooperates, $\partial p_{t+1}/\partial \rho > 0$ implies that $\partial T\left(k_{t+1}|k_t\right)/\partial \rho$ will be positive for states k_{t+1} that have more cooperators. In addition cooperation is greatly enhanced when a PEF has a strong imitative component. If agent's don't directly look at their neighbors actions when updating their strategy, then it will move agents in the direction of defecting more often.

4. The sum of the negative effect on the PEF of neighbors induced by an "invader" is less than the positive effect from rival's cooperation on the invader. As mentioned above, each time an agent moves he lowers the probability of cooperation for the agents to which he is newly connected. And, as these agents defect more often their neighbors become more likely to defect more often. Calculating the sum of this effect is difficult because the effect propagates to neighbors of neighbors, etc. Thus the network structure needs to be able to cut off this cascade in an efficient manner so that the negative sum is small compared to the positive effect on the invader from interacting with cooperators. But the network structure also needs to allow cooperative behavior to spread efficiently. Interestingly the lattice structure appears most efficient at balancing these two effects. As we've seen above neighborhood size and neighborhood network structure can make a difference. In the circle model, larger neighborhood sizes are important, as is the case with the lattice model.

7 Conclusion

This paper has explored the effects of movement on the attainment of cooperation in a spatial Prisoner's Dilemma using both analytical techniques and agent-based simulations. Movement has the potential to either allow cooperating agents to avoid defectors or to allow defectors to invade cooperators. We provide a series of computational models that help to explain when each result is likely to be obtained.

First we show that when agents evolve their mixed-strategies based on their actions and their rivals' actions, these rules give rise to a Markov chain over all possible states of the "cooperative system." Next we focus on one particular probability evolution function (PEF), with a strong imitation effect. This PEF has two absorbing states: all agents cooperate and all agents defect. We then look at the evolution of cooperation in three network structures: a circle, a lattice and discrete neighborhoods. We look at two sets of outcomes. The first set is when agents do not move. The second set is when agents have the opportunity to move to a new (randomly chosen) location based on a comparison of payoffs with their current neighbors to the payoffs they would get if they moved into the new neighborhood.

The simulations show that network structure is an important determinant of the evolution of cooperation and that network structure can affect whether agent movement improves or worsens cooperation. In the circular network, we see that movement strictly increases cooperation. Also we see that neighborhood size matters, with smaller and larger networks fostering cooperation more than medium size neighborhoods. For the lattice networks, movement can dramatically increase cooperation, especially when the neighborhood size is 8 rather than 4. Larger neighborhoods are more effective at propagating cooperation. Finally, with the discrete networks we see that movement decreases cooperation because there is no mechanism to prevent the spread of defecting agents within a neighborhood.

This work has aimed to explore how the inclusion of "conscious" movement in an RPD framework can affect the evolution of cooperation. While we have found some strong results with our simulations, this work is a first attempt at looking at this relatively simple addition to the RPD. Future work can explore this game using more analytical methods. In addition, it would be interesting to see how agent heterogeneity–in terms of agents' probability evolution functions–affects the results of the game, with and without movement.

References

[1] Aktipis, C. Athena. 2004. Know When to Walk Away: Contingent Movement and the Evolution of Cooperation. *Journal of Theoretical Biology*, 231: 249-260.

[2] Axelrod, Robert. 1984. *The Evolution of Cooperation.* New York: Basic Books.

[3] Brauchli, K, T. Killingback, and M. Doebeli. 1999. Evolution of Cooperation in Spatially Structured Polulations. *Journal of Theoretical Biology*, 200, 405-417.

[4] Dugatin, L.A. 1992. The Evolution of the Con-Artist. *Ethology and Sociobiology*, 13, 3-18.

[5] Dugatin. L.A. and D.S. Wilson. 1991. ROVER: A Strategy for Exploiting Cooperators in a Patchy Environment. *American Naturalist*, 138, 687-701.

[6] Enquist, M. and O. Leimar. 1993. The Evolution of Cooperation in Mobile Organisms. *Animal Behavior*, 45, 747-757.

[7] Ferriere, R. and R.E. Michod. 1995. Invading Wave of Cooperation in a Spatial Iterated Prisoner's Dilemma. *Preceedings: Biological Science*, 259, 77-83.

[8] Fudenberg, Drew and Eric Maskin. 1991. Evolution and Cooperation in Noisy Repeated Games *American Economic Review*,Papers and Proceedings 80(2), 274-279. Cambridge: The MIT Press.

[9] Fudenberg, Drew and Jean Tirole. 1991. *Game Theory.* Cambridge: The MIT Press.

[10] Hales, David. 2001. Tag-based Cooperation in Artificial Societies. Ph.D. Thesis, University of Essex.

[11] Hanaki, Nobuyuki, Alexander Peterhansl, Peter S. Dodds, and Duncan J. Watts. 2007. Cooperation in Evolving Social Networks. *Management Science*, 53(7), 1036-1050.

[12] Imhof, Lorens A., Drew Fudenberg, Martin A. Nowak and Robert M. May (2005). *Proceedings of the National Academy of Sciences of the United States of America*, 102(31), 10797-10800.

[13] Janssen, Marco A. 2008. Evolution of Cooperation in a One-shot Prisoner's Dilemma Based on Recognition of Trustworthy and Untrustworthy Agents. *Journal of Economic Behavior and Organization*, 65(3-4): 458-471.

[14] Killingbach, T. and M. Doebeli. 1996. Spatial Evolutionary Game Theory: Hawks and Doves Revisited. *Proceedings: Biological Sciences*, 263, 1135-1144.

[15] Nowak, Martin A. and Robert M. May. 1992. Evolutionary Games and Spatial Chaos. *Nature*, 359: 826-829.

[16] Nowak, Martin A. and Karl Sigmund. 1998. Evolution of Indirect Reciprocity by Image Scoring. *Nature*, 393: 573-577.

[17] Nowak, Martin A., Sebastian Bonhoeffer, and Robert M. May. 1994. Spatial Games and the Maintenance of Cooperation. *Proceedings of the National Academy of Science*, 91: 4877-4881.

[18] Poundstone, William. 1992. *Prisoner's Dilemma*, Doubleday, NY, NY.

[19] Riolo, Rick L. 1997. The Effects of Tag Mediated Selection of Partners in Evolving Populations Playing the Iterated Prisoner's Dilemma. Santa Fe Institute Working Paper # 97-02-016.

[20] Riolo, Rick L., Michael Cohen, and Robert Axelrod. 2001. Evolution of Cooperation without Reciprocity. *Nature*, 414: 441-443.

[21] Schweitzer, Frank. Laxmidhar Behera and Heinz Mühlenbein. 2002. Evolution of Cooperation in a Spatial Prisoners Dilemma. *Advances in Complex Systems*, 5: 269-299.

[22] Schluessler, R. 1989. Exit Threats and Cooperation Under Anonymity. *Journal of Conflict Resolution*, 33, 728-749.

[23] Vanberg, V. J. and R.D. Congelton. 1992. Rationality, Morality and Exit. *American Political Science Review*, 86, 418-431.

[24] Wilhite, Allen. 2006 Economic Activity on Fixed Networks. Handbook of Computational Economics, Volume 2, Eds. Leight Tesfatsion and Kenneth L. Jud. North-Holland