Endogenous Movement and Equilibrium Selection in Spatial Coordination Games

Troy Tassier[1], Jason Barr[2], David Hagmann[1]

[1]Department of Economics, Fordham University, Bronx, NY 10458, USA
{tassier, hagmann}@fordham.edu
[2]Department of Economics, Rutgers University, Newark, NJ 07102, USA
jmbarr@rutgers.edu

Abstract. We study the effects of agent movement on equilibrium selection in spatial coordination games with Pareto dominant and risk dominant Nash equilibria. We incorporate agent movement, and thereby partner selection, into the strategies of agents in a class of spatial coordination games. Our primary interest is in understanding how endogenous partner selection on networks influences equilibrium selection in games with multiple equilibria. We use agent based models and best response behaviors of agents to study our questions of interest. In general, we find that allowing agents to move greatly increases the probability of selecting the Pareto dominant Nash equilibrium in coordination games.

Keywords: Coordination Games, Movement, Equilibrium Selection

1 Introduction

Agent movement has previously been studied in spatial prisoner's dilemma games. In this field, researchers are interested in whether the ability of agents to move favors the invasion of defecting agents into neighborhoods of cooperators, or whether the ability of agents to move allows cooperators to escape defectors. Previous research [1], [2] suggests that the ability of agents to move enhances rates of cooperation on average. For example, Aktipis [1] studies the behavior of a "walk-away" strategy in a spatial prisoner's dilemma game where agents cooperate if a rival cooperated in a previous period or move to a new location if the rival defected in the previous period. She finds that walk-away is a successful strategy when placed in an Axelrod [3] style tournament among commonly studied strategies such as tit-for-tat. Barr and Tassier [2] study the rates of cooperation and evolution of mixed strategies in a spatial prisoner's dilemma game where agents are allowed to move. They find that the opportunity to move greatly enhances the probability of agent cooperation across many (but not all) network structures.

In this paper we study how the ability of agents to move affects equilibrium selection in spatial coordination games with Pareto dominant and risk dominant Nash equilibria. Consider the following generic two agent, two strategy simultaneous game:

		Player 2	
		X	Y
Player 1	X	a , a	b , c
	Y	c , b	d , d

Throughout the paper we assume $a>c$, $d>b$ such that there exist two pure strategy Nash Equilibria, X,X and Y,Y. Agents attempt to coordinate with their play partners on one of the two Nash equilibria. Thus our game of interest is a standard 2x2 coordination game. Further, we assume that $a>d$ such that X,X is the *Pareto dominant Nash equilibrium.* Harsanyi and Selton [4] define equilibrium Y,Y to be a r*isk dominant Nash equilibrium* if $(a-c)(a-c) < (d-b)(d-b)$ which is equivalent to $a+b < c+d$. Our primary interest in this paper will be with payoffs assigned such that Y,Y qualifies as risk dominant. We study equilibrium selection in these games using an agent-based model.

Note that there is a tension between agents attempting to coordinate on the Pareto dominant versus the risk dominant Nash equilibrium. All agents would prefer to coordinate on X,X because each agent receives a larger payoff than in Y, Y. But, should coordination not occur (one agent playing X and the other playing Y) the agent attempting to coordinate on X, X is penalized with a low payoff of b. More importantly, as b decreases, playing X becomes more risky and playing Y becomes more attractive.

There exists a large literature on the long run selection of equilibria in these games without agent movement. As examples, Ellison [5], Kandori, Mailath, and Rob [6], and Young [7], study equilibrium selection in an evolutionary framework where agents are randomly matched with game partners. They find that the risk dominant Nash equilibrium is the unique stochastically stable equilibrium when agents have a small probability of making mistakes in strategy selection. Morris [8] studies the spread of a Pareto dominant Nash equilibrium in spatial games where agents play a coordination game on various topologies. He finds that a Pareto dominant equilibrium may be favored in some network based coordination games if the number of neighbors in the network expands at an intermediate rate (quickly, but not too quickly.)

In this paper, we explore how endogenous agent movement (which has not been previously studied in coordination games) affects the equilibrium selection results described above. Specifically, agents are located on a two-dimensional lattice. Agents play a coordination game with each nearest neighbor on the network in each period. Agents choose a best response to last periodȏs play by their neighbors as their action in the current period. Using agent based modeling, we study the evolution of the agent strategies and the attainment of Pareto vs. risk dominant Nash equilibria. We find that the Pareto dominant Nash equilibrium is much more likely to be attained when agents have the ability to move on the network and choose game play partners than when agents are not allowed to move.

2 Our Model

Each *run* of our model proceeds as follows. In the initialization procedure, N agents are created and assigned a random location on an LxL lattice with fixed boundaries (not a torus).[1] Note that we set N to be less than LxL so that there are some vacant locations. In addition, each agent is assigned a random strategy, X or Y, according to a specified probability distribution. This random initial strategy is played in the first period of the model.

Following the initialization procedure, each agent plays the coordination game described in the previous section, with each neighbor on the lattice. In each period, each agent chooses a single action that must be played with every neighbor. Agents choose this action as follows: The action in period one is assigned by the initialization procedure. In each successive period agents choose an action that is a best response to their neighborsô actions in the previous period. Specifically, an agent calculates her average payoff for both strategy X and strategy Y against the strategies chosen by each of her neighbors in the previous period. Whichever strategy yields a larger average payoff is chosen in the current period. Ties are broken by the agent playing the strategy that she most recently played.[2]

Following agent game play, each agent is individually given an opportunity to move to a new vacant location on the lattice with probability m. If the agent is given this opportunity, a random vacant location is chosen from among the set of vacant locations on the lattice with uniform probability. The agent then calculates the best response strategy at the new location, X or Y, and the corresponding average payoff of that strategy. The agent then compares the average payoff of the best response at the new location to the average payoff of the best response at the current location. If this average payoff is greater at the new location she moves there. Otherwise she remains at her current location. We then repeat this game play procedure until we generate equilibrium behavior.[3] We repeat the entire process (initialization and game

---

[1] This assumption is made without loss of generality. We also have experimented with a torus topology and not found qualitatively different results.

[2] Note that we make the strategy selection using average payoffs because, when movement is allowed, two locations may have different numbers of neighbors (due to some locations being vacant and also due to edge effects on the lattice.)

[3] Note that because we are using strict best response for strategy selection, some equilibria contain period to period switching by agents. As a simple example, if there are two agents whose only neighbor is the other it is possible that they never coordinate and continuously switch strategies in each period. For instance one agent may choose X, Y, X, etcé while the other chooses Y, X, Y, etcé With strict best response there is no way to break this cycle. Because these cycles continue in perpetuity, we consider them to be equilibrium behavior.

play) for R runs for a specified set of parameters. We take averages over these R runs and report results for each parameter set below.

3 Results

We report average results below for $R=200$ runs of each parameter set. We are primarily interested in equilibrium selection differences when movement is allowed in the model versus when movement is not allowed. Therefore we vary the probability of a movement opportunity between $m=0$ and $m=1$ across different sets of runs and compare the equilibrium selection results.

3.1 Variation in Risk

We begin this comparison with the following payoff selections: $a=2.0$, $c=0.0$, $d=1.0$ so that we have the following normal form game representation:

		Player 2	
		X	Y
Player 1	X	2 , 2	b , 0
	Y	0 , b	1 , 1

We then vary b at intervals between 0 and -6. Note that $b<1$ implies that Y, Y is a Nash equilibrium and $b<-1$ implies that Y, Y is a risk dominant Nash equilibrium. Initially we set $N=100$ and $L=12$, so that there are 144 locations: 100 locations with agents and 44 vacant locations. For these runs we also set the probability of playing X in the first period equal to 50%. Thus, on average, there will be 50% of agents playing X in period one and 50% playing Y in period one.

In Table 1 we report the average percent of agents that play strategy X in equilibrium when no movement is allowed, $m=0$, and when movement is allowed for each agent in every period, $m=1$, as a function of the payoff parameter b.

To begin, note the behavior of the model when no movement is allowed. As expected, the percent of agents playing the Pareto dominant Nash equilibrium strategy, X, decreases as the payoff b decreases. (As b decreases playing X becomes more risky.) Further, recall that when $b<-1$, Y, Y becomes a risk dominant Nash equilibrium. And, we see in the table that for $b>-1$ a majority of agents play the Pareto dominant Nash equilibrium strategy. But, for $b<-1$, a majority of agents play strategy Y that corresponds to the risk dominant Nash equilibrium. This corresponds well with the existing literature on equilibrium selection in coordination games reported in the introduction. Although our model is different (network based matching versus random matching) the risk dominant equilibrium is still favored in our model as predicted.

Table 1. Percent of agents that coordinate on the Pareto dominant strategy, X, as a function of the payoff b (average over 200 runs of the model.) Other payoffs: a=2.0, c=0.0, d=1.0. No movement, m=0 vs. movement, m=1. 12x12 lattice.

Payoff b	m=0	m=1
0	98.6%	100%
-0.5	82.0%	100%
-1.0	50.4%	100%
-1.1	23.5%	98.7%
-1.5	18.4%	94.2%
-2.0	6.7%	60.9%
-2.5	2.1%	42.6%
-3.0	1.3%	16.0%
-4.0	0.3%	6.5%
-4.5	0.2%	4.1%
-5.0	0.2%	2.7%
-6.0	0.1%	1.8%

More interesting is the comparison of the results when movement is not allowed to the results when movement is allowed. As one can see in the table, allowing movement greatly increases the probability that the Pareto dominant strategy is played in equilibrium. Further it is still possible to generate small numbers of agents who play the Pareto dominant strategy in equilibrium even when doing so is very risky, when b is very small.

3.2 Variation in Population Size

Next we consider how the density of agents on the lattice changes equilibrium selection. Here we leave the payoffs for a, c, and d as above and set b=-1.5 (an intermediate value for b where we had strong effects for movement.) Again we set the initial strategy distribution equal to 50% for each strategy.

As reported in Table 2, changing the agent density does little to the equilibrium selection results when movement is not allowed. However, when movement is allowed, a more dense lattice makes it more difficult to coordinate on the Pareto dominant Nash equilibrium. This occurs for two reasons: First, when the network is very dense, it may be more difficult to find a vacant location near a small pocket of X, X coordinating agents. Second, because the network is more dense it is difficult for X, X coordination to spread. Each location will have more occupied locations adjacent to it. Thus it may be difficult for a small group of agents to ñtipò toward the Pareto dominant Nash equilibrium. When the network is less densely populated it may be easy to find small groups of unconnected agents that can coordinate on the X, X Nash equilibrium. Then once these agents coordinate, movement into adjacent vacant cells can allow this equilibrium to spread outward. This process is more difficult when the lattice is densely populated.

Table 2. Percent of agents that coordinate on the Pareto dominant strategy, X, as a function of population size, N (averages over 200 runs of the model). Payoffs: a=2.0, b=-1.5, c=0.0, d=1.0. No movement, m=0 vs. movement, m=1. 12x12 lattice.

Number of Agents	m=0	m=1
70	22.0%	98.6%
85	20.5%	98.7%
100	18.4%	94.2%
115	18.2%	79.5%
130	18.0%	62.5%

3.3 Variation in Initial Strategy

Finally, we vary the initial strategy distribution in the population. Again we leave the payoffs unchanged with a=2.0, b=-1.5, c=0.0, and d=1.0. But, we vary the initial percentage of agents playing strategy X in period one from 20% to 70% (with the complement playing Y). We report the results in Table 3.

Table 3: Percent of agents that coordinate on the Pareto dominant strategy, X, as a function of the initial percentage of strategy X agents (averages over 200 runs of the model.) Payoffs: a=2.0, b=-1.5, c=0.0, d=1.0. No movement, m=0 vs. movement, m=1. 12x12 lattice.

Percent of Initial Agents Playing X	m=0	m=1
20%	0.2%	4.3%
30%	1.0%	21.5%
40%	4.7%	71.0%
50%	18.4%	94.2%
60%	49.2%	100%
70%	81.0%	100%

There are two items of note in the table. First, the initial distribution of agents has a large effect on the equilibrium selected. Moving the initial percentage of agents playing X slightly above (below) 50% greatly increases (decreases) the probability of agents coordinating on the X, X Nash equilibrium. Second, movement can act with large force to counteract the initial distribution. For instance, when only 40% of agents play strategy X in the initial period, only 4.7% of agents play X in equilibrium when movement is not allowed. But, allowing movement increases the equilibrium incidence of X to 71%. Even when only 30% of agents play strategy X initially, movement can allow a significant percentage of agents to coordinate on the Pareto efficient strategy. In this case 21.5% of agents play X in equilibrium. But to some extent this understates the effect of movement. In 26 of these 200 runs, all 100 agents coordinate on the Pareto dominant Nash equilibrium. Thus, as stated above, the ability to move and select game play partners can have a large effect on the equilibrium selected. More importantly this effect moves the population strongly toward selecting the Pareto dominant Nash equilibrium more frequently.

4 Discussion

The main results of this paper find that movement greatly improves ability of agents to coordinate on a Pareto superior equilibrium in coordination games. The results should be viewed with three respects in mind: First, the results here directly extend the work of others on the effects of movement in Prisonerōs Dilemma games, such as Aktipis [1] and Barr and Tassier [2], to other games. Second, the results are an extension to the literature on equilibrium selection in coordination games cited in the introduction. Movement can greatly increase the ability of agents to avoid the risk dominant Nash equilibrium in favor of the Pareto superior Nash equilibrium. Third and most generally, the results of movement in these games should also be viewed as an example of endogenous partner selection. Overall, the results reported in this paper suggest that, when agents are able to choose game play partners, superior game outcomes can be expected. These superior outcomes may be cooperation in a Prisonerōs Dilemma or coordination on superior although more risky Nash equilibria.

5 Future Work

We are currently working on a variety of extensions to the results reported above. We mention a few of them here. Most prominently we are investigating additional models of strategy selection by agents. In the results reported above we only consider agents that choose a strategy that is a best response to last periodōs play by neighbors. There are a number of additional strategy selection models of interest to us. For instance agents may respond to more than one periodōs past play. This would introduce a greater degree of persistence in agentsô strategy choices and may increase or decrease the propensity to coordinate on a given equilibrium. There also are many alternative behaviors to best response dynamics such as stimulus response behaviors or imitative dynamics which may produce interesting equilibrium selection results. In addition, investigation of an evolutionary model which selects among a class of agents behaviors, such as those described above, would be of interest.

All agents in the model above are homogeneous in terms of risk preference (all agents have linear preferences over payoffs) and payoffs. Particularly in the case where agents have the ability to move, the introduction of heterogeneous agents would be interesting. For instance, would agents with different risk preferences and the ability to move be able to better coordinate? Would these agents sort themselves spatially according to preferences? Similarly, suppose that a small set of agents had a very large (small) set of payoffs associated with strategy X, payoff a and b above. Would this preference for strategy X in a small set of agents have a large or a small effect on the equilibrium selection results? Would the effect be seen across the entire population of agents or would these agents separate themselves into a small pocket on the lattice and have no effect on the other agents? Many question such as these would be interesting to explore.

Finally, we report results for a two-dimensional lattice above. While we feel these results are of general interest, we also are performing series of runs with other topologies and are investigating how variations in the topology of the network change the equilibrium selection outcomes.

References

1. Aktipis, C.A. (2004). Know when to walk away: Contingent movement and the evolution of cooperation. *Journal of Theoretical Biology*, 231 (2), 249-260.
2. Barr, J. and Tassier, T. (2010). Endogenous neighborhood selection and the attainment of cooperation in a spatial prisoner's dilemma game. *Computational Economics*, 35, 211-234.
3. Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.
4. Harsanyi, J.C. and Selten, R. (1988). *A General Theory of Equilibrium Selection in Games*, Cambridge, MA, MIT Press.
5. Ellison, G. (1993), Learning, Local Interaction, and Coordination. *Econometrica*, 61, 1047-1071.
6. Kandori, M., Mailath, G. J., and Rob, R. (1993). Learning, mutation, and long run equilibria in games. *Econometrica*, 61, 29-56.
7. Young, P. (1993). The Evolution of Conventions. *Econometrica*, 61, 57-84.
8. Morris, S. (2000). Contagion. *The Review of Economics Studies*, 67, 57-78.