PD games on networks

by

Allen Wilhite

Department of Economics and Finance
University of Alabama in Huntsville
Huntsville, Al 35763
wilhitea@uah.edu

prepared for CEF 2005

abstract

We tend to interact with same people, day after day. Might this affect our behavior? In an abstract fashion, we look at this question. To model this repeated interaction with a small subset of the entire population we place agents on the nodes of a network and have them play a prisoners' dilemma game exclusively with their neighbors. We then alter the payoffs of the game and the topology of the network to see if, when, and to what degree cooperation survives. We find widely divergent aggregate decisions across networks and across payoffs. But, there is commonality as well. It seems clear that some networks, or some organizational structures, are more conducive to fostering cooperation

**This study is in its infancy; this manuscript and results preliminary. We urge caution when considering our results and we welcome comments and suggestions.*

PD Games on Networks

I. Introduction

Unexpected things happen when a prisoners' dilemma game is repeatedly played on a network. For example, Nowak and May (1993, 1994) place players on a grid and find that the distribution of strategies can be chaotic, making it impossible to predict the actions of any individual player. Eshel, Samuelson and Shaked (1998) look at players on a ring and even though the defection strategy dominants any individual pairing, they find that cooperation not only survives but spreads to dominate the network. Wilhite (2005) compares the distribution of strategies on a series of networks and finds that each network leads to a unique long term distribution of cooperation and defection.

But network topology affects more than the eventual number of cooperators and defectors. Given a particular network, one set of payoffs can lead to one dynamic outcome while slightly different payoffs may create a completely different set of decisions by the actors. And, the long run distribution of strategies may be stable with one set of payoffs but cycle with a different payoff matrix. This paper investigates PD games played on networks and the distribution of strategies selected by players. Within the remarkable variation that arises, we find order. For example, all regular networks have ranges of payoffs that trigger identical dynamics but, the internal dynamics change at certain critical values. Moreover, these critical values are the same for *all* regular-size networks of a particular order.

II. The game, playing the game, and the playing field:

the game:

Two person symmetric games can be expressed in their normal form with the familiar 2X2 matrix.

	C	D
C	(a,a)	(b,c)
D	(c,b)	(d,d)

Symmetric prisoners' dilemma games refer to those in which $c > a > d > b$ and $a+d > c+b$ (although this second condition is not pertinent in the following applications). A simplified version of this game illustrates the activity that provided the motivation for this paper. For now, let's restrict three of the parameters such that $a=1; b=d=0$. While this isn't strictly a prisoner's dilemma game (because $b = d$ thus the dual-defection result is only a weak Nash equilibrium), it still demonstrates the behavior of interest when $c>1$ and allows us to analyze the game by adjusting this single parameter, c. In section IV we will remove these simplifications to investigate true PD games.

play

This is an evolutionary game in which strategies evolve over time with repeated play. However, agents don't learn in the sense that they remember and react to the play of other specific agents. Instead, agents initially play a randomly selected strategy (C, cooperate or D, defect), playing the same strategy with all of their neighbors. After one round of play, they collect their payoffs and observe the play and payoffs of each neighbor. Preceding the next round of play, each agent updates his strategy by imitating the most successful strategy adopted in his neighborhood. Notice that in each round each agent plays the same strategy with all of his neighbors, that is, if agent *i* has chosen the defection strategy, he defects with all of his neighbors. Other flavors of imitation strategies have also appeared in the literature such as imitating the most successful *average* strategy (Eshel, *et.al.*,1998) or the most popular strategy (Ellison and Fudenberg, 1993) but we have selected one of the simplest, agents imitate their most successful neighbor.

networks (the playing field)

Each agent plays this PD game with a select group of other agents defined by the network. In this paper we concentrate on networks in which each node is connected to k nodes, k being a constant. We call these networks with regular-size neighborhoods. For instance, if $k = 4$, then each agent is connected to four neighbors.[1]

We introduce our study by exploring three simple networks in which $k = 4$, the ring, the grid, and a portion of a tree. To visualize a *k4* ring think of a network in which agents are spread around a circle being connected to four others, the two agents on each side. Regardless of which agent is selected his neighborhood looks like every other agent's neighborhood. The grid also sees agents laid out as on a chessboard each agent connected to the four closest agents, one above, one below and one on each side. Boundaries are eliminated by wrapping the edges around to connect creating a torus. Again, select any agent on this network and his neighbors line up in the same configuration. Constructing a tree within which each agent has four neighbors is impossible because agents out on the tips of the branches have only one neighbor, but for everyone else, the neighborhood is constant. For example, start with a central agent and connect him to four agents. Each of these branch off to connect to another three agents, each of which branches thrice, and so forth. The interior of this tree consists of agents with identical neighborhoods. But at the tips of the branches many agents have only one neighbor, the immediately higher neighbor. To avoid that boundary issue, we create a large tree and concentrate on the interior.

III. Results of some simple experiments:

Identical games played on different networks behave differently. To illustrate these differences consider the eventual distribution of strategies on these three networks as we adjust the defection payoff, *c*. Table 1 shows the final distribution of strategies adopted by a network of 1600 agents engaged in a repeated PD game following the rules outlined above, starting with a randomly selected distribution of strategies in which

[1] A regular-size neighborhood is a bit more inclusive than a network of constant degree. Constant degree networks (or more explicitly constant degree graphs) are those in which each node has the same number of edges. Regular size (a constant number of neighbors) includes all constant degree networks but also allows for self-play, i.e. an agent can be a neighbor with himself or more importantly engage in self-play. Self play allows us to think of the "decision maker" occupying a node as an organization (family or firm) that also plays the game internally. In addition self-play allows us to consider networks with a specific topology and an odd number of neighbors. For example, without self-play you cannot have a regular-sized ring network with an odd number of neighbors.

approximately 10% initially defect. The final tallies are an average of 20 separate runs for each network and payoff combination.

Table 1
Eventual distribution of strategies and dynamics
of PD games on three networks

	1 < c < 4/3	4/3 < c < 3/2	3/2 < c < 2	2 < c < 3
Ring	3-cycle: 87% cooperators	3-cycle: 85% cooperators	stable: 52% cooperators	all defectors
Tree	stable: 73% cooperators	some 2-cycles: 20% cooperators	2-cycles: 42% cooperators	some 2 cycles: 12% cooperators
Grid	stable/2-cycle: 68% cooperators	chaotic: 35% cooperators	long cycle: 22% cooperators	all defectors

Table 1 shows how the dynamics of play lump into groups. Reading across the table, we see that there are ranges of payoffs that yield identical results. Consider a ring. If the defection payoff, *c*, is between 1 and 4/3, the network converges to a population in which about 87% on the individuals cooperate, while the remaining defectors roll through a three-period cycle. It does not matter what the initial value of *c* is, as long as it lies in the identified range. In fact, given the same initial distribution of strategies, the long run distribution cooperators and defectors is identical for any value of *c* as long as it lies between 1 and 4/3. If $c > 4/3$, the dynamics abruptly change and we see fewer cooperators and, if *c* is even higher such that it crosses the next threshold ($c > 3/2$), the dynamics of play shifts into a third pattern. These initial experiments suggest that there is a phase transition at these critical values of *c*. Table 1 also shows that for these three networks, these critical values of *c* are identical for each network; the phase transitions are triggered by the same payoff parameters. In the next section we will see this applies to any regular-sized network.

Second, while each network shares phase-transition points, they display different behavior within a specific parameter range. For example, suppose c = 1.4, reading down the center column we see that a ring topology leads to a population consisting of 85% cooperators. That same population distributed on a tree yields a population in which only 20% of the population cooperates and a grid yields an average level of cooperation of about 35%.

Third, the nature of each network's long-term steady-state distribution of strategies differs. Again reading down the center column, a ring evolves into a 3-period cycle of cooperation and defection, a tree is usually stable but occasionally evolves a 2-period cycle and the grid yields a chaotic pattern of cooperation and defection that on average yields 35% cooperation, but leaves the specific strategy of any particular agent uncertain. These dynamic characteristics also depend on the topology of the network that defines the players.

This paper investigates the origins of these three characteristics: (i) what are the critical payoffs that trigger the transition from one set of results to another on a particular network, (ii) why do different networks display different strategy patterns given a particular set of payoffs, and (iii) why do some networks converge to a stable distribution of strategies while others fall into a cycle with some agents perpetually switching their strategy. Our first task is to formally derive the critical values that identify these phase transition points. We approach this by investigating the circumstances under which a particular strategy will spread. In this formal derivation we will abandon the simplified payoffs and explore the general form, symmetric, PD game. We also allow a network to have any size neighborhood, maintaining the restriction that that size is constant through the network. When examples are useful, we will return to our *k4* applications.

IV. Spreading strategies

1. critical values and phase transition

The eventual distribution of strategies depends on the circumstances under which a particular strategy spreads to its neighbors. So, when does defection spread? Consider two neighboring agents, *i* and *j* where agent *i* is currently defecting and agent *j* is cooperating. Agent *j* switches to defection if his most successful neighbor is currently defecting and is earning more than himself. He retains his cooperative strategy if his most successful neighbor (including himself) is cooperating. Focusing on agent *j*'s decision, suppose agent *i* is his most successful defecting neighbor and agent *j** is his most successful cooperative neighbor (for the remainder of this paper, defecting agents will have superscripts *i* and cooperative agents will be *j*'s). Let n^i be the number of agent *i*'s neighbors and n_C^i the number of *i*'s *cooperating* neighbors. Note, agents *i*, *j*, and *j** are all in *j*'s neighborhood, agents *i* and *j* are in *i*'s neighborhood, but agent *j** may or may not be a neighbor of *i*. Also note that agent *j* could be that most successful agent, agent *j**. With this notation, agent *j* switches to defection if agent *i* earns more than *j**. Given the payoffs of the prisoners' dilemma game agent *j* switches to defection if equation (1) holds.

$$(1) \qquad c(n_C^i) + d(n^i - n_C^i) > a(n_C^{j*}) + b(n^{j*} - n_C^{j*})$$

In the simple game $a = 1, d = b = 0$, this becomes $c(n_C^i) > n_C^{j*}$.

Rearranging (1) agent *j* switches to defection if

$$(2) \qquad c > \frac{a(n_C^{j*}) + b(n^{j*} - n_C^{j*}) - d(n^i - n_C^i)}{n_C^i}.$$

By definition of a prisoner's dilemma game, this inequality always holds if $n_C^i \geq n_C^{j*}$ (because $c > a$ and $d > b$). Thus the potentially interesting cases are those in which agent *j* has a neighbor with more cooperating neighbors than agent *i*, or if $n^{j*}{}_C >$

n^i_C. In those cases, defection spreads if c exceeds the critical values given in the left half of Table 2 (assuming a *k4* network).

Table 2
Critical values for a *k4* network

	Defection spreads: (2) true; (4) false		Cooperation spreads: (2) false (4) true	
	neighborhood configuration	payoffs (PD)	neighborhood configuration	payoffs (PD)
1	$n_C^i = 1;\ n_C^{j*} = 4$	$c > 4a\text{-}3d$	*	*
2	$n_C^i = 1;\ n_C^{j*} = 3$	$c > 3a + b - 3d$	$n_C^{i*} = 1;\ n_C^j = 3$	$c < 3a + b - 3d$
3	$n_C^i = 1;\ n_C^{j*} = 2$	$c > 2a + 2b - 3d$	$n_C^{i*} = 1;\ n_C^j = 2$	$c < 2a + 2b - 3d$
4	$n_C^i = 2;\ n_C^{j*} = 4$	$c > 2a - d$	*	*
5	$n_C^i = 2;\ n_C^{j*} = 3$	$c > \frac{3}{2}a + \frac{1}{2}b - d$	$n_C^{i*} = 2;\ n_C^j = 3$	$c < \frac{3}{2}a + \frac{1}{2}b - d$
6	$n_C^i = 3;\ n_C^{j*} = 4$	$c > \frac{4}{3}a - \frac{1}{3}d$	*	*
7	$n_C^i \geq n_C^{j*}$	$c > a$ (any PD game)	$n_C^{i*} \geq n_C^j$	$c < a$ (no PD game)

* cells are empty because by definition $n_C^j < 4$

So, if a particular defector has two cooperating neighbors ($n_C^i = 2$) and one of those has another neighbor with three cooperating neighbors ($n_C^{j*} = 3$) defection spreads only if c exceeds the payoffs in row 5 ($c > \frac{3}{2}a + \frac{1}{2}b - d$). If c is less than that expression defection no longer spreads.

On the other hand, even if c falls below that value, cooperation may not spread either. For cooperation to spread, a defecting agent's most successful cooperating neighbor must earn more than his most successful defecting neighbor. Labeling agent *i*'s most successful defecting neighbor as *i**, and his most successful cooperating neighbor as *j*, agent *i* switches to cooperation if

$$c(n_C^{i*}) + d(n^{i*} - n_C^{i*}) < a(n_C^j) + b(n^j - n_C^j) \qquad (3)$$

or

$$c < \frac{a(n_C^j) + b(n^j - n_C^j) - d(n^{i*} - n_C^{i*})}{n_C^{i*}} \qquad (4)$$

As for the case of spreading defection we can use equation (4) to calculate the critical values of the payoffs that trigger a phase transition in the spread of cooperation. For a *k4* network those calculations are in the right hand side of Table 2.

The key to the dynamic change across networks is that these two criteria, equations (2) and (4), are not symmetric. In total they can involve four different neighborhoods, those defined by agents, *i*, *i**, *j*, and *j**. But the spread of defection depends on *i* and *j** while the spread of cooperation depends on *i** and *j*. Given a particular set of payoffs, equations (2) and (4) show that it is the number of cooperating neighbors that determines the relative earnings of agents and determines their decision to switch strategies or maintain their current play. By construction $n_C^{i*} \geq n_C^{i}$ and $n_C^{j*} \geq n_C^{j}$, but, n_C^{i*} and n_C^{i} can be greater or less than n_C^{j*} and n_C^{j}. Consequently, for any specific pair of neighbors, *i~j*, inequality (2) can be true or false, inequality (4) can be true or false, and they can both be false, but both inequalities cannot be true. These combinations give rise to the spread or contraction of a particular strategy. For example, if (2) is true and (4) is false, defection spreads from *i* to *j*, if (4) is true and (2) is false, cooperation spreads from *j* to *i*, and if both are false, both agents retain their current strategy.

A more intuitive way to present the combinations of payoffs that trigger the transition from one distribution of strategies to another is to view those payoffs spatially. Setting *a = 1* and holding *b = 0* for a moment, the combinations of payoffs that yield identical distributions of strategies can be graphed as in Figure 1 by letting the vertical axis be the payoff value for *c* and the horizontal axis the payoff value for *d*.

Figure 1
Critical payoff combinations in *k4* networks

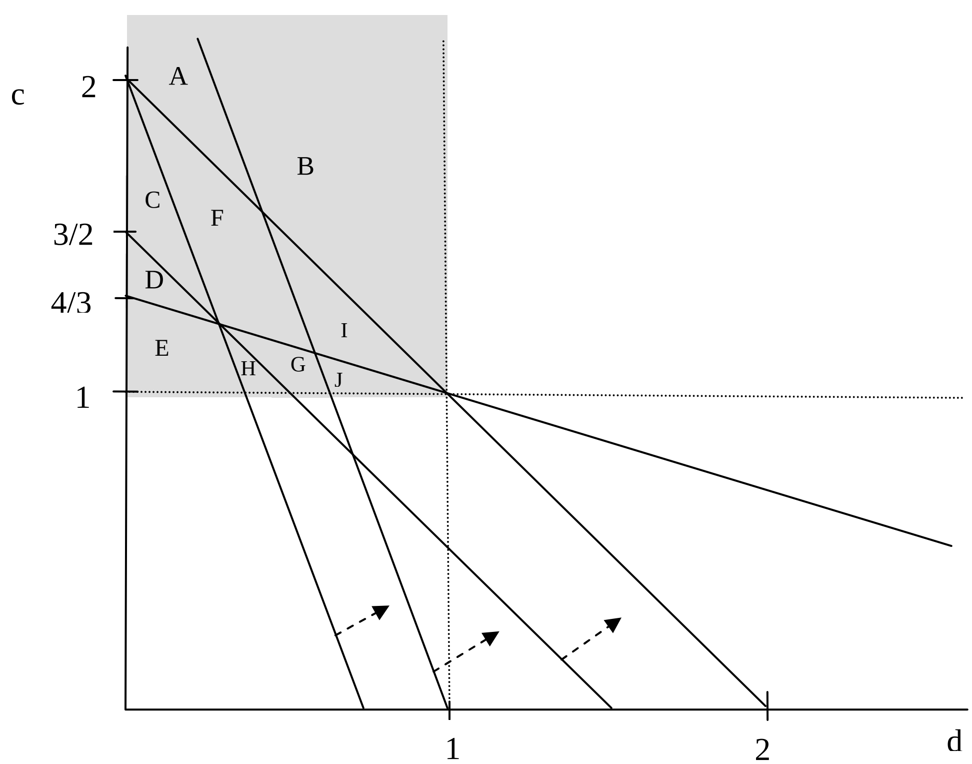

PD games occupy the shaded region of Figure 1 ($c > a = 1 > d > b = 0$). Instead of the four parameter ranges given in Table 1 (when only *c* changed) there are now ten regions of payoffs that trigger unique dynamics (A through J).[2] As before, any combination of payoffs that lie within a particular two-dimensional region leads to identical dynamics on any specific network. Moving from one region to another, say by increasing *d*, triggers a phase transition to different dynamics and distribution of strategies. Finally, Figure 1 can also show how a change in the sucker's payoff, *b*, affects the game. The dashed arrows in the bottom of Figure 1 show how the three indicated boundaries shift outward as *b* increases. This further alters the specific combinations of payoffs that yield identical results, that is, this changes the shape of the regions a bit. In addition three new regions can emerge as *b* rises, but as before, the dynamics within each of these spaces is identical for any initial distribution of strategies.

2. Differences across Networks:

Just as the inequalities (2) and (4) can be used to identify combinations of parameters that trigger different dynamics in the same network, they can also be used to see how different networks respond to the same payoffs. Given a set of payoffs, the deciding parameters in the inequalities (2) and (4) are the size of the neighborhood and the number of neighbors who cooperate. Different networks define different neighborhoods and different overlapping of neighbors. Consider a section of a *k4* ring in which agents have neighborhoods that overlap with either three or four common neighbors as shown in the following diagram.

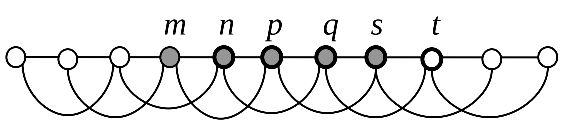

Agent *p*'s neighborhood consists of the shaded agents while agent *q*'s neighborhood is indicated with the bold nodes. Neighbors *n*, *p*, *q*, and *s*, are common to both neighborhoods. Contrast that with following piece of a *k4* grid.

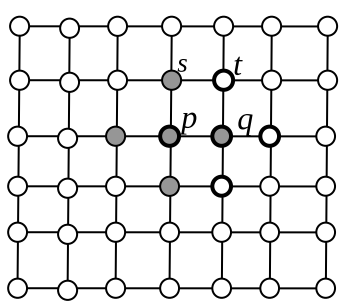

In the grid, each pair of neighbors has no additional common neighbors. Agent *p*, whose neighborhood consists of the shaded agents, overlaps with agent *q*'s neighborhood (bold nodes) only through agents *p* and *q*. However, while *p* and *q* have no common neighbors, their neighbors are neighbors. For example agents *p* and *q*'s northern

[2] Note, if d=0, the tick marks on the vertical axis lie at the critical values of c presented in Table 1.

neighbors (agents *s* and *t*) are neighbors. This "once removed" overlapping is absent in other networks, such as a tree. As you can see in the tree below, *p* and *q* share no neighbors nor are their neighbors neighbors, nor do they share neighbors.

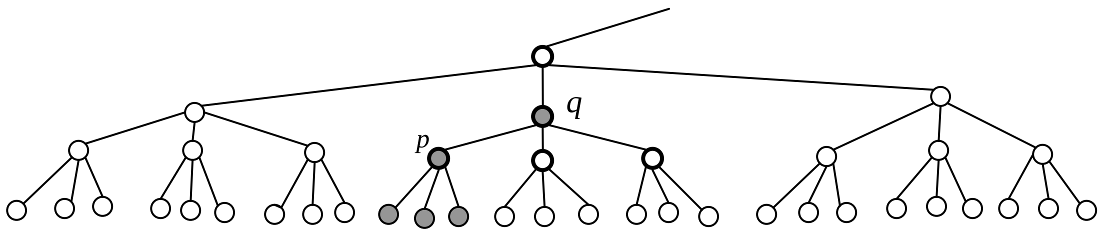

The overlap of neighbors' neighborhoods is critical to the spread of a strategy. Suppose agent *i* is defecting. Spreading depends on the number of cooperators and defectors in agent *i*'s neighborhood relative to the number of cooperators and defectors in agent *j**'s neighborhood (recall that *j** is the most successful cooperating agent in *j*'s neighborhood and $i \sim j$). If the neighborhoods overlap a great deal (as in the ring) when one agent has many cooperators in his neighborhood, it is likely that his neighbors have cooperating neighbors as well (they share so many neighbors.) Visa versa, if *j* has many defecting neighbors (reducing his payoff) it is likely that agent *i* will have many defectors in his neighborhood which also reduces his payoff. But in a network in which neighborhoods overlap little, the payoff to agent *i*'s strategy is much more independent of the strategies of *j*'s and *j**'s neighbors.

Consequently, identical payoffs lead to different aggregate behavior in different networks. Consider a network in which every agent in a single neighborhood is playing the defection strategy and everyone else is cooperating. Focus on agent i, $i \sim j$, i defecting j, cooperating. In a grid, $n^{i*}_C = n^i_C = 3$, while $n^j = 3$ and $n^{j*} = 4$. Holding $a = 1$, defection spreads if $c > 4/3 - 1/3\, d$. We are now in regions A, B, C, D, F, or I (depending on the values of b and d) in Figure 1.

For a contrast, suppose this defecting neighborhood is on a ring. Now $n^{i*}_C = 2$ and n^i_C = 1 or 2 while n^{j*} = 4 and n^j = 2 or 3. Holding $a = 1$, defection spreads if $c > 2 - d$; we are in the regions labeled A and B in figure 1. Thus, if $d < 1$ (which is true by the definition of a PD game) there exists a values of c for which defection spreads on the grid but does not spread on the ring. In this example, any combinations that land us in regions C, D, F, or I will trigger a different response in the ring versus the grid. So, while all regular networks share the same critical values, the dynamics that arise for a given set of payoffs differs across networks.

3. Cycles vs stable outcomes

The final network characteristic to be explored is the presence of cycles in some of the long run distribution of strategies. Stable or unchanging long run outcomes are easily understood. Some pockets of agents playing one strategy or the other are either sufficiently resilient to change that they maintain the status quo, or some pockets are sufficiently weak that they are driven to extinction. In either case a stable distribution of strategies emerges.

Cycles typically involve decisions made in a single neighborhood. Consider agents *i~j* when *i* is defecting and *j* is cooperating. Recall, the spread of either defection or cooperation depends on the number of cooperators in their neighborhood (n_C^i, n_C^{i*}, n_C^j, *and* n_C^{j*}). Suppose the payoffs in a particular network are such that cooperation is spreading. Once again, focus on a single neighborhood consisting of defectors and surrounded by defectors. As cooperation spreads this set of defectors diminishes. It diminishes to a single neighborhood of defectors, and finally to a single agent playing the defection strategy. We now have a sole defector surrounded by cooperators. For this agent, $n^{i*}_C = n^i_C = 4$ and $n^{j*}_C = 4$. In the next round of play, everyone in his neighborhood copies the defection strategy (because $c > a$). But now we are once again in a pattern in which cooperation is spreading and the number of defectors shrinks once again. It dwindles to a single defector who then converts his neighborhood and so forth; a cycle.

In a *k4* ring 3 period cycles emerge if the payoffs lie in regions D, E, or H, for example, if $c = 1.4, a = 1, d = .2, b = 0,$ we get a three period cycle as shown below (the shaded nodes represent defectors).

Figure 2
Three-period cycle in a RING

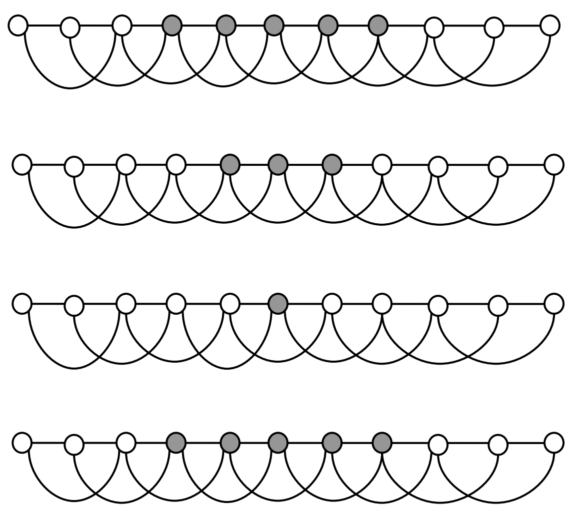

There is another three-period cycle on the ring that arises when spreading cooperation squeezes defectors down to a group of two neighbors. In region D, these two defectors are earning enough to spread to their entire neighborhood (six agents). Now, however, the string of six agents begins to unravel at the ends to four then again two agents. This creates a three-period cycle [2, 6, 4], [2, 6, 4],…

Two-period cycles arise in *k4* tree networks emerge when the proper mix of cooperation and defection meets at a particular junction. For instance, if the payoffs lay in regions C, D, F, or I and agents have adopted the strategies as shown in Figure 3 (shaded nodes are defectors), the tree oscillates in a 2-period cycle.[3]

Figure 3
Two-period cycle in a TREE

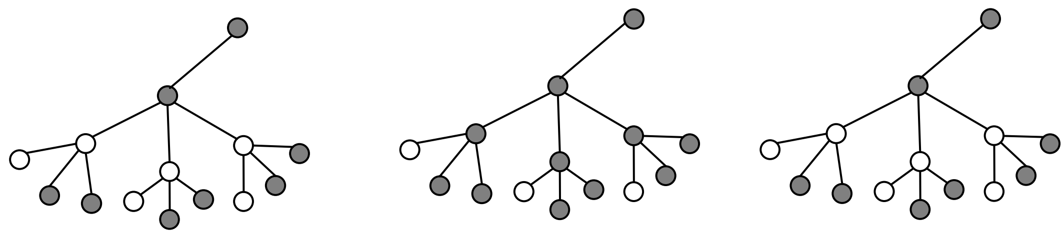

V. Extensions:

While this manuscript looks at PD games, the critical payoff regions shown in Figure 1 can be used to examine other games as well. First, instead of considering action C as cooperation and D as defection just consider them as some general choice. That is, an agent can take action C or action D. As we have seen, PD games are those whose payoffs appear in the upper left-hand quadrant of Figure 1 (the shaded region). The lower left-hand quadrant consists of the stag hunt (when $c > d$) and games of coordination in which the CC choice is the payoff dominant equilibrium (when $1 > d > c$). The lower right-hand quadrant consists of games of coordination in which the DD selection is payoff dominant (when $d > 1 > c$). As with PD games, the aggregate number of agents who choose action C or action D differs as we move from region to region.[4] And, as before, given a particular set of payoffs, aggregate coordination depends on the topology of the underlying network.[5]

The upper right-hand quadrant of Figure 1 is another game, one with a single dominant equilibrium strategy shared by both players. In this region the payoffs are either, $d > c > a > b = 0$ or $c > d > a > b = 0$ and both players choose action D. Notice that this region is not subdivided into smaller areas, the dynamics and eventual distribution is the same; all agents choose action D.

Finally, the examples used in this manuscript concentrated on a few *k4* networks, but these results generalize to all regular-sized, uniform networks. Consider, for example, playing a PD game on a *k6* network. The dynamics and eventual distribution of

---

[3] In Figure 3 it is assumed that the tree extends further in every direction, and the hidden generations have adopted the strategy of the last neighbor showing. That is, a defector on the pictured "rim" has additional neighbors who are defecting and a cooperator on the rim has cooperating neighbors.

[4] Actually the lines appearing in Figure 1 includes only those tradeoffs that trigger phase transitions in prisoners' dilemma games. There are additional regions to consider with games of coordination (and some of these regions are unimportant) but inequalities (2) and (4) generate will all of the pertinent combinations.

[5] For an extensive look at coordination games on networks, particularly grids, see Young (2002, 1993) and Morris (2000).

strategies of *k6* networks depends on the payoffs and the network that defines the players. Once again each network has its own equilibrium set of strategies that change abruptly when the payoffs cross a threshold. And, within a particular payoff areas, different networks generate different behavior. The transition points for any size network can be calculated using equations (2) and (4). Figure 4 shows the payoffs again assuming $a = 1$ and $b = 0$ for a *k6* network. As before, prisoners' dilemma games lie in the shaded region, but there are now 24 parameter combinations that create different dynamics across networks.

VI. Conclusions:

In many economic and social situations the individuals with whom we interact are the same people, time after time. We shop at the same stores, work with the same group of colleagues and our homes are surrounded by a largely unchanging set of neighbors. Networks are a method of formalizing these relationships. With this in mind this manuscript looks at a particular type of interaction, the prisoners' dilemma, to find that that the aggregate behavior of individuals is markedly affected by the topology of their network. Viewing networks as alternative organizational structures, is seems quite clear that the level of cooperation is affected by the type of organization in which one is situated. There seem to be certain features shared across networks, the critical payoffs that trigger the different types of aggregate behavior for instance, but there are stark differences as well. Some organizational structures seems to be much more conducive to cooperative behavior, other structures seem to be hostile to its survival and still other networks see perpetual changes as agents cycle through periods of cooperation versus defection.

References:

Eshel, I., L. Samuelson, and A. Shaked (1998), "Altruists, egoists, and hooligans in a local interaction model." *The American Economic Review*, Vol. 88, 1:157-179.

Ellison, G. and D. Fudenberg (1993), "Rules of Thumb for Social Learning." *Journal of Political Economy, Vol. 101, 4:612-643.*

Morris, S. (2000), "Contagion," *Review of Economic Studies*, 67: 57-78.

Nowak, M. and R. May (1993), *International Journal of Bifurcation and Chaos*, 3, 35-78.

Nowak, M, and R. May (1994), *International Journal of Bifurcation and Chaos*, 4, 33-56.

Wilhite, A. (2005). "Economic activity on fixed networks.", *Handbook of Computational Economics: Vol II.* ed. L.Tesfaction and K. Judd. (North-Holland: Amsterdam).

Young, H. P. (1993), "The Evolution of Conventions," *Econometrica*, Vol. 61, 1:57-84.

Young, H.P. (2000), "The Diffusion of Innovations in Social Networks." working paper, Department of Economics, John Hopkins University; Center on Social and Economic Dynamics, the Brookings Institute; and The Santa Fe Institute.

Figure 4
Critical payoff combinations for
prisoners' dilemma games played on *k6* networks