Cooperation and Friendship Choice Heuristics in Dynamic Local Interaction Networks

Trevor G. Dobson[1]
University of Virginia

Abstract

I built an agent-based model that simulates economic interaction in a social network. In the model, agents dynamically break and create network links based on whether their neighbors cooperate in a repeated Prisoner's Dilemma game. Agents use heuristics to predict which potential neighbors will yield the highest payouts and form links with those neighbors, but some heuristics which are individually beneficial have significant negative externalities. These heuristics can change the network's structure, increasing its vulnerability to systemic defection cascades which reduce payouts for all agents.

1 Trevor Dobson is an alumnus of the University of Virginia, where he studied Economics, Biology and Computer Science. This research was completed as a senior thesis for the Economics Distinguished Majors Program. Trevor is currently employed as a Financial Analyst in the Algorithmic Trading department of Goldman Sachs. Comments are welcome at TrevorGDobson@gmail.com.

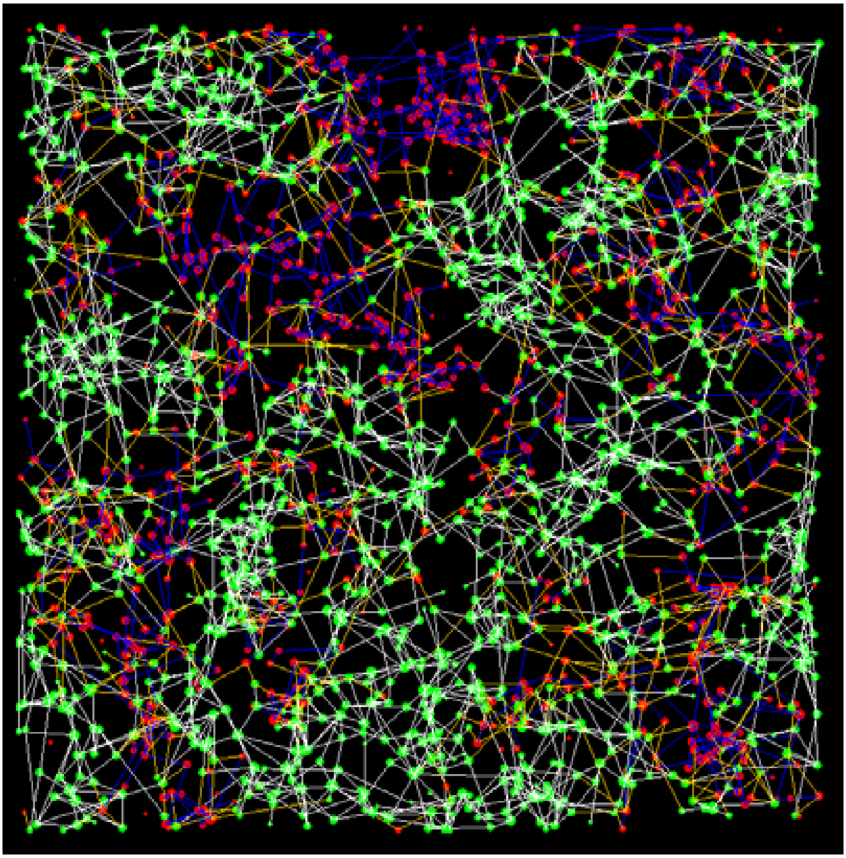

Simulation with 2000 agents playing the Locality strategy. Green agents are cooperating; red agents are defecting. Agents' size is proportional to the number of neighbors they possess. Links between pairs of cooperators are white, links between pairs of defectors are blue, and links between cooperators and defectors are orange.

I. Introduction

The global economy is a complex, interconnected network of markets. Crises that begin as isolated, local events spread rapidly through the network in unpredictable ways. Loss of confidence in one market not only depresses asset prices locally, but spurs further loss of confidence in neighboring markets. Modern financial crises propagate in ways traditionally associated with epidemiology and infectious disease. Financial contagion spreads across nations, as in the Asian financial crisis of 1997 and 1998, and across asset classes, as in the current financial crisis, which began in sub-prime mortgages but expanded to the markets for commercial debt and equity.

I study this behavior by abstracting investment decisions and physical space. In my model, agents play a repeated Prisoner's Dilemma game with their neighbors in a local interaction social network, which represents the repeated interaction that occurs in any economy. Agents use a simple imitation strategy to choose their behavior in each round—either cooperation or defection—and attempt to improve their payoffs by selectively breaking and creating links with other agents. In some situations, the network can fall apart as links are repeatedly broken and agents choose not to cooperate. This type of defection cascades across the network and is analogous to financial contagion.

Because agents break links rationally, they end relationships with partners who do not cooperate in the prisoner's dilemma game in order to form relationships with potentially more cooperative agents. Given that agents observe their neighbors' behavior, they have the information necessary for rational link destruction.

Rational link creation is more difficult. Because agents receive higher payouts from links with cooperative agents, they have an incentive to preferentially link with cooperating agents. I depart from standard equilibrium theory by assuming that agents lack complete knowledge of other agents' behaviors, and only know the recent behavior of their immediate neighbors. Agents use limited information heuristics to approximate whether potential new neighbors would cooperate in the future.

In the short run, several of the heuristics increase agents' average payouts but have a negative externality. This externality is caused by structural changes in the network when induced by certain heuristics used by a critical mass of agents. These structural changes can cause network-wide defection for specific periods of time, reducing the payoffs to all agents. It is this emergent phenomenon that I seek to explore with this simulation.

II. Literature Review

The Prisoner's Dilemma and Cooperation

In the Prisoner's Dilemma, actors choose to either cooperate or defect. Cooperation benefits both players while defection benefits only the defector at the expense of the cooperator. There are many examples of the Prisoner's Dilemma in biology and economics, including the classic anecdote involving two criminals who are offered plea bargains by the police if they testify against their partner. In this example, the police have separately detained two suspects for a crime. Both suspects are given an offer: if they testify against their partner, they will receive a lighter sentence. If neither testifies, they are both convicted of small offenses, and receive 3 years in jail. If one testifies and the other does not, the testifier is granted immunity (o years in jail), but the other is convicted and given 10 years in jail. If both testify, they are both convicted but receive lighter sentences: 7 years in jail.

The setup of the prisoner's dilemma is always the same. Players face four possible payoffs, for which I use Axelrod's terminology.[2] In Axelrod's construct, if both players cooperate, they receive a reward payoff, R. If both defect, they receive a penalty, P. Each player has reason to defect, because if they defect while the other player cooperates, they receive a "tempting" payoff, T, while the other player receives the "sucker's" payoff, S. The payoffs appear as follows:

Player 1 / Player 2	**Defect (D_2)**	**Cooperate (C_2)**
Defect (D)	(P, P)	(T, S)
Cooperate (C)	(S, T)	(R, R)

Table 1. Theoretical Prisoner's Dilemma Payoffs

For the prisoner's dilemma condition to hold, it must be the case that $T > R > P > S$. If this is the case, then defecting is a strictly dominant strategy ($T > R$, $P > S$), though both players prefer mutual cooperation to mutual defection ($R > P$). The final condition of the game is that both players prefer mutual cooperation to alternating between unilateral cooperation and unilateral defection ($R > (T+S / 2)$). These conditions hold in the previous example, because actors always have an incentive to testify (7 years < 10 years, 0 years < 3 years), but both agents are better off if neither testifies (3 years < 7 years).

In a simple, one-round Prisoner's Dilemma game, defection is the domi-

2 Robert Axelrod and William Donald Hamilton. The Evolution of Cooperation. Science 211, no. 4489 (1981): 1390.

nant strategy, and mutual defection is the only Nash Equilibrium. This contradicts common sense and empirical observation, as cooperation occurs between rational actors in human interaction and between irrational actors in nature. So why does cooperation exist? The simplest explanation of cooperation in human society is coercion. Institutions like governments and religious and trade organizations can change payoffs by reducing the temptation to cheat (T), so that mutual cooperation becomes a Nash Equilibrium.

While coercion by third parties can explain some cooperation, mutual cooperation occurs in many situations without coercion, i.e. those situations where the net payoffs in the Prisoner's Dilemma game obey the original rule, $T > R > P > S$. One explanation for cooperation in the absence of coercion is the use of potential future payoffs in repeated games as rewards and punishments for behaviors. In indefinitely repeated games, cooperation can be justified as a selfish strategy if both players expect the other player to play a "grim trigger" or "tit for tat" strategy. In these cases, the expected punishment given to cheaters, in the form of non-cooperation by their counterparts, is worse than the expected gain from defection, which makes long-term cooperation a Nash Equilibrium in repeated games.[3]

Another way to explain cooperation in Prisoner's Dilemma situations is to relax the assumption that agents play rationally. In imitation models, agents make decisions by imitating other agents they perceive as successful, rather than calculating best responses to opponents' anticipated strategies. This type of behavior is rational when agents lack the capacity to forecast the dynamic strategies used by other agents in their environments. In constructs with this type of imitation, agents can sustain supra-Nash payoffs and cooperate indefinitely when cooperators imitate each other.

Eshel, Samuelson and Shaked give an excellent example of this type of model in their 1998 work, "Altruists, Egoists, and Hooligans in a Local Interaction Model."[4] In this model, agents choose between three behaviors—altruism, egoism and hooliganism—in which they respectively contribute to a public good, act as free riders, or actively detract from the public good for their own benefit. The model shows that a cluster of altruists imitating one another can be robust to invasion by egoists and hooligans. The model yields persistent contribution to the public good, despite the fact that non-cooperation is a dominant strategy in a strictly rational sense. Theodore Bergstrom's review, "Evolution of Social Behavior: Individual and Group Selection," is a good

3 Robert Axelrod and William Donald Hamilton, Ibid.

4 Ilan Eshel, Larry Samuelson, and Avner Shaked. "Altruists, Egoists, and Hooligans in a Local Interaction Model." The American Economic Review 88 no. 1 (1998).

summary of the earlier analytical work done in this field.[5]

Imitation strategies make local interaction networks susceptible to rapid transitions between majority cooperation and majority non-cooperation states. Stephen Morris's paper "Contagion," published in 2000, shows that information spread by imitative strategies can move rapidly through a local interaction network, with dynamics analogous to those of infectious disease.[6] This phenomenon is robust to the introduction of various forms of noise and randomness.[7] When these defection cascades occur, the payouts for all agents in the social network drop from the cooperative level to the mutually non-cooperative level. For this reason, the structures of interaction networks have a direct effect on the payouts agents receive.

Social Networks

My work models Prisoner's Dilemma games in social networks. Social networks have always been an important factor in economic interactions, and modern technologies have opened up new avenues for research. The modern study of networks includes the study of the "small world phenomenon."[8] Small world networks are unique; their elements are highly clustered, but are also rapidly traversable and searchable.

Clustering implies that nodes which share mutual neighbors are more likely to be neighbors themselves. The rapid traversability of a network is usually measured with a parameter called the network diameter, which is the average minimum number of network links required to move from a random node in the network to another random node.

Random networks have high traversability but low clustering. As random networks increase in size, the network diameter increases at a logarithmic rate. Local, or ordered, networks have the opposite property; agents in a specific area are more likely to be connected, but the network diameter is much larger.

Watts and Strogatz demonstrated that, by randomly replacing a small proportion of the links in a highly ordered local network, they could generate hybrid networks that share the properties of both types. They called these hybrids "small world networks," because the random links allow rapid traversal from one end of the network to the other, lowering network diameter without sacrificing clustering.

5 Theodore Bergstrom. "Evolution of Social Behavior: Individual and Group Selection." Journal of Economic Perspectives 16 no.2 (2002):67-88.

6 Stephen Morris. "Contagion." Review of Economic Studies 67 (2000): 57-78.

7 In Ho Lee and Akos Valentinyi. "Noisy Contagion Without Mutation." Review of Economic Studies 67 (2007):47-56.

8 D. Watts and S. Strogatz. Collective dynamics of 'small-world' networks. Nature 393 no. 684 (1998): 409-10.

Dynamic Social Networks

In "Cooperation in Evolving Social Networks," Hanaki, Peterhansl, Dodds, and Watts theorize a network in which agents play a repeated Prisoner's Dilemma game with multiple neighbors in a local interaction network.[9] Agents have a fractional chance of being allowed to end or begin relationships with other agents in each round. Hanaki et al. study the rational decisions made by agents when determining which strategies to play against new partners.

The model developed by Hanaki et al. begins by assuming that new partners are selected randomly from the general pool of agents. They also theorize a variant of their model in which agents preferentially form links with "friends of friends," which they call a "triadic closure bias." In theory, the triadic closure bias is justified by the fact that cooperative agents continually end relationships with non-cooperators, so cooperative agents' friends must be more cooperative than the population average. So, if a given agent's friends are more likely to be cooperative, then the friends of that given agent's friends are also more likely to be cooperative.

The triadic closure bias is a simple example of what I call a "friendship choice heuristic." Friendship choice heuristics are rules of thumb used to preferentially create links with cooperating agents by reasonably predicting other agents' behavior with only limited information. I study the properties of five friendship choice heuristics, including their immediate effects on the payouts of individual agents and their indirect effects on the payouts of all the other agents in the network. The five heuristics I examine are:

1. **Random:** Agents playing the Random strategy select new neighbors randomly from the pool of all agents; all other agents have an equal probability of being selected. This is my baseline heuristic.

2. **Referral:** Agents playing the Referral strategy select new friends from a pool of "friends of friends." The more mutual friends a potential neighbor has with the selecting agent, the more likely that potential neighbor is to be selected. This is akin to requiring referrals or recommendations for new business partners or friends, and assigning preference to the number of recommendations received.

3. **Popularity:** The probability of selecting a given agent is proportional to the number of links that potential neighbor currently has. For example, an agent with four friends is twice as likely to be selected as an agent with only two friends.

9 N. Hanaki, A. Peterhansl, P. Dodds, D. Watts. "Cooperation in Evolving Social Networks." Management Science 53 no. 7 (2007):1036-50

4. Unpopularity: This is similar to the Popularity strategy, but the relationship between a node's number of links and likelihood of being picked is reversed. For example, a node with four links is less likely to be picked than a node with one link.

5. Locality: This heuristic assumes that nodes are fixed in two-dimensional space, and gives probability weights to each node inverse to their distance from the selecting node. For the purposes of this heuristic, agents are assigned random locations in a one unit by one unit square.

The following diagram shows how each of these five heuristics would work:

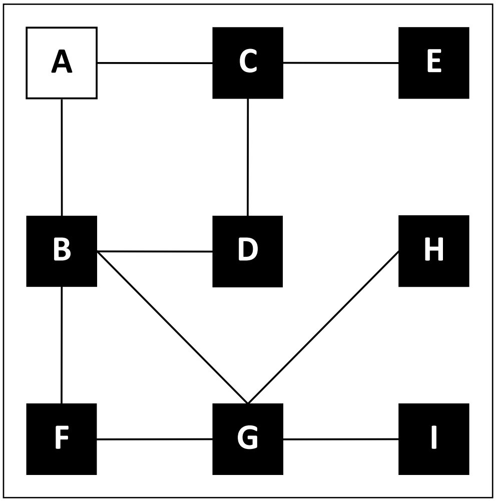

Figure 1. Link Creation Example

Different heuristics favor different types of nodes. In this case, A is creating a new link, and D, E, F, G, H and I are all candidates. Under each of the friendship heuristic regimes, they would have the following relative probabilities of being selected:

Heuristic	Pr(D)	Pr(E)	Pr(F)	Pr(G)	Pr(H)	Pr(I)
Random	1 (16.6%)	1 (16.6%)	1 (16.6%)	1 (16.6%)	1 (16.6%)	1 (16.6%)
Referral	2 (40%)	1 (20%)	1 (20%)	1 (20%)	0 (0%)	0 (0%)
Popularity	2 (18.2%)	1 (9.1%)	2 (18.2%)	4 (36.4%)	1 (9.1%)	1 (9.1%)
Unpopularity	2 (11.8%)	4 (23.5%)	2 (11.8%)	1 (5.9%)	4 (23.5%)	4 (23.5%)
Locality	1.4 (11.1%)	2 (15.7%)	2 (15.7%)	2.2 (17.6%)	2.2 (17.6%)	2.8 (22.3%)

Table 2. Link Creation Example Relative Probabilities

III. Methodology

Initialization

The simulation begins by creating an array of agents and linking each of them to four other agents. The simulation runs for 5000 rounds before it starts recording payout and cascade data, to guard against any biases from the random distribution of links and behaviors when the system is initialized.

Iteration of Rounds

Each round of the game consists of four sections. First, agents decide whether to break or create new links with other agents. Second, agents decide which behavior they will adopt for the current round by imitating their most successful neighbor. Third, agents play a Prisoner's Dilemma game with their neighbors. Finally, the program updates running average payoffs and displays results as necessary.

Updating Links

Each agent's decision of whether to break links is simple. Because agents expect their neighbors to continue to play as they have in the past, agents should break links with any neighbors who have defected in the previous round. Agents should break these links regardless of whether they themselves defect, because even if the neighbor simultaneously defects, the agent is better off finding a new partner who may be a cooperator. Each agent has a 20% chance of being allowed to break a link in each round. This represents the inherent difficulty in ending long-term relationships.

Agents create new links if and only if they have just broken links. This is the only circumstance in which new links are created. Only the agent that breaks the link can create a new link, meaning that any time a link is broken, the defecting agent loses a neighbor, and some other agent gains a neighbor. Thus, the agent breaking and creating links has no net change. This keeps the number of links in the simulation constant, and also prevents any agent from intentionally changing the number of links it possesses. This creates an incentive for agents to cooperate, because if they defect, they will be ostracized.

As mentioned previously, agents have a variety of options for selecting new partners with whom to form links. There are five options: the Random strategy, the Popularity strategy, the Referral strategy, the Unpopularity strategy and the Locality strategy.

Agents playing Random select new partners from the general pool of agents using the Java.Math library's random number generator. If one of the agent's preexisting friends is selected, the algorithm selects again.

Agents playing Referral link to agents with whom they share a common neighbor but are not currently linked. Agents implement the strategy by iterating through their neighbors and making a list of link candidates (their neighbors' neighbors). Because they may share multiple mutual neighbors with certain candidates, those candidates may be added to the list multiple times, increasing the chance that they will be chosen. A random agent from the list of candidates is chosen to be a new neighbor.

The Popularity strategy is more complex than the Random and Referral strategies. A simple implementation would be to use a lottery system like the Referral algorithm, but to give each candidate a number of lottery entries according to the number of links they possess rather than the number of links they receive. While straightforward, this method would be computationally intensive, forcing the simulation to iterate through the entire set of agents every time Popularity was used. To increase the efficiency of the simulation, a different algorithm is used that yields the same Popularity-weighted distribution but uses less processing power.

Candidates are randomly selected from the pool of all agents. Selected candidates are then assigned a score based on their popularity divided by a constant. The simulation then generates a second random number, and if the candidate's score is greater than the random number, the candidate becomes a new neighbor of the original agent playing the Popularity strategy. If the candidate's score is lower than the random number, then a new random candidate is selected from the pool of all agents, and the process is repeated until a satisfactory candidate is found. Each individual's probability of being selected is still directly proportional to his or her number of neighbors, because the probability that an individual is selected to become a candidate is a constant, and the probability that a candidate's score will be greater than the random number is proportional to the number of neighbors that candidate has.

The Unpopularity and Locality strategies operate by the same mechanism but use different scores. In the Unpopularity strategy, scores are inversely related to the number of neighbors the candidate has, such that a candidate with two neighbors would have a score of ½, while a candidate with four neighbors would have a score of only ¼. In the Locality strategy, the score is the inverse of the candidate's distance from the selecting agent, so that candidates closer to the agent are given a higher chance of being selected.

Behavior Selection (Noisy Imitation)

The third portion of each round is behavior selection. Agents decide what decision to play based on their observations of their immediate neighbors. Agents iterate through their set of neighbors and select the neighbor with the highest payoff in the previous round. They then change their behavior to match that neighbor's behavior in the previous round. In this way, agents choose their behavior by imitating their most successful neighbor in the previous round.

The imitation process is noisy. Each round, there is a 1 in 1000 chance that agents will "innovate," randomly selecting their behavior in the next round instead of imitating their neighbors. In some cases, this noise can initiate movements between homogenous cooperation and homogenous defection, but in most cases it has little effect on the long run behaviors of agents.

This procedure allows agents to make reasonable responses to their environment even though they lack the capacity to calculate rational best responses. Agents' imitative responses can resemble rational behavior in many situations. For example, an agent whose neighbor defects may imitate that neighbor in the next round, effectively punishing that neighbor by defecting, which resembles the appropriate response dictated by the tit for tat strategy. In other situations, agents may continue to cooperate despite defection by some neighbors if their most successful neighbors are cooperators. Agents who continue to cooperate despite their neighbor's defection can respond by severing the link with the defecting neighbor, as an alternative form of punishment. Generally, imitative responses are similar to best responses.

Payout Generation

The game itself is a Prisoner's Dilemma with the following payoffs:

A \ B	Defect	Cooperate
Defect	0.1, 0.1	1.5, 0
Cooperate	0, 1.5	1, 1

Table 3. Simulation Prisoner's Dilemma Payoffs

This version of the Prisoner's Dilemma game obeys the standard rules, T > R, P>S, R > P and R > (T+S)/2. Agents play this game against each of their neighbors. Their payoff for the round is the sum of the payoffs from each of their interactions. It should be noted that this game specification is arbitrary, and it affects the outcome of the simulation. A simulation in which agents are judged based on the average payoff they receive from their interactions would yield different results, as would a simulation where agents play a public goods game rather than a Prisoner's Dilemma.

Reporting Statistics

Agents' payoffs are stored in vectors, making it easy to find the network's average payoff at any given time or over the life of the simulation, or to find those averages for the set of agents using a particular heuristic.

The reporting process is complicated slightly by the technical difficulty of storing every agent's payoff history in limited memory. To work around this problem, the simulation periodically takes snapshots of the relevant simulation averages and wipes the individual payoff histories (excepting the most recent rounds, which are used in agents' behavior decisions the following turn).

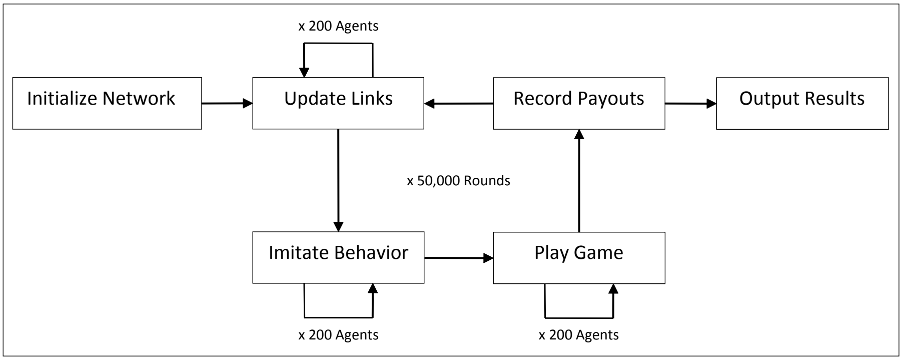

Figure 2: Simulation Summary

The model is implemented in Java version 1.6. For the purposes of my analysis, I ran the simulation for 55,000 rounds at a time, varying the parameters between runs. I omitted the first 5,000 rounds from the data, to avoid initialization effects. I then recorded the average payoffs for all agents over all rounds of each run of the simulation, as well as separate averages for the payoffs for agents using each of the heuristic types.

I also recorded the number and total length of the defection cascades that occurred in each run and calculated the average length of each cascade. I defined cascades according to the following criterion: a cascade begins whenever the fraction of cooperating agents in the network falls below 5% and ends when the fraction of cooperating agents returns to 40% or higher. It should be noted that these thresholds do not include all agents in the simulation; as mentioned before, islands (agents with no links) are unconnected to the network and therefore unaffected by defection cascades. My results are not sensitive to the precise details of how cascades are defined.

Using this methodology, I ran a series of experiments in which a fraction of the agent population use the Random friend choice heuristic and the remainder use a different strategy. In each experiment, I ran several hundred instances of the simulation, distributed over the range of possible combinations of the two strategies. To accelerate this Monte Carlo process, I ran the simulation on the University of Virginia's Centurion cluster, a group of 32 IBM eServer 325s with Dual 1.6 GHz Opteron252 processors and 2GBs of RAM each.

IV. Results

The first result I observed when running my simulation was qualitative. Certain combinations of parameters yield defection cascades in which a predominantly cooperative network transitions to being predominantly non-cooperative in only a few (usually 4-5) rounds. The spread of defection through the network in these situations is rapid and complete, similar to the spread of disease in an epidemiological model. After several hundred rounds, the system recovers, reverting to cooperation over the course of 10-20 rounds.

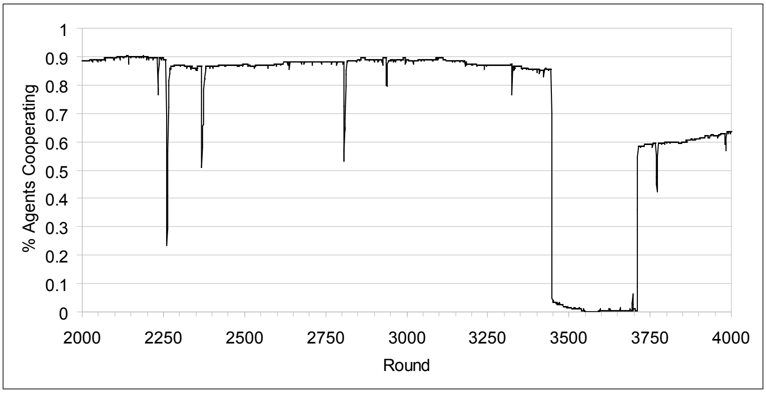

Chart 1: Cascade Example

Chart 1 shows the fraction of agents cooperating in a typical simulation run over 2000 rounds. In the example, a severe defection cascade occurs at round 3445, which causes every agent connected to the network to defect over the course of the next four rounds. Islands, nodes that are unconnected to the network, are unaffected by these cascades, so there is some residual cooperation until round 3500. After about 200 more rounds, the system recovers, and the cascade is over. Unlike the minor defection events at rounds 2259, 2369 and 2807, the major cascade affected all of the agents in the network and as a result lasted much longer and had a more significant effect on the agents' average payouts over the life of the simulation. In addition, during the cascade many agents had all of their links cut, thereby removing them from the network, so that when the network recovered shortly before round 3750, only 60% of the agents returned to cooperation.

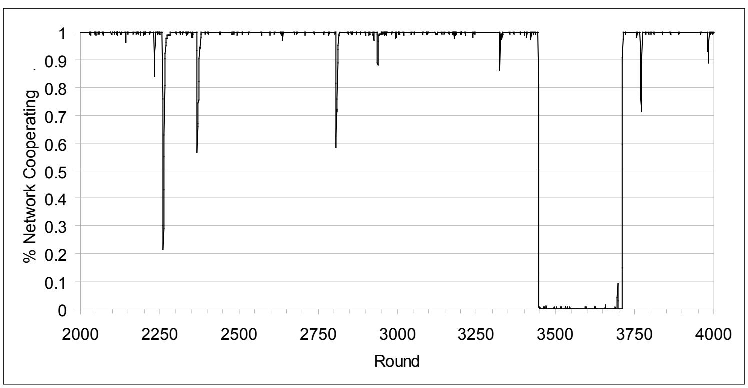

Chart 2: Cascade Example without Islands

When island nodes are excluded from the data, the cascades look more extreme. Every agent defects at the beginning of a cascade, and every agent returns to cooperation at the end.

Random vs. Popularity

The Random and Popularity strategies were the first pair of heuristics I compared. The data show that when the majority of agents play Random, the agents playing Popularity consistently perform better. At the same time, when more than half of the population plays the Popularity strategy, first the Random and then both types of agents' payoffs are diminished due to a large increase in the severity and number of defection cascades. This is born out in the negative correlation between the proportion of agents playing the Popularity strategy and average payouts and in the positive correlation between the proportion of agents playing the Popularity strategy and the frequency and severity of defection cascades.

In addition to this negative trend, there is a structural break in the data when roughly 40% of agents use the Random strategy. For this reason, my regressions of the data use a dummy variable to distinguish between data from above and below this structural break, so that OLS can be used to specify a piece-wise linear regression for the data.

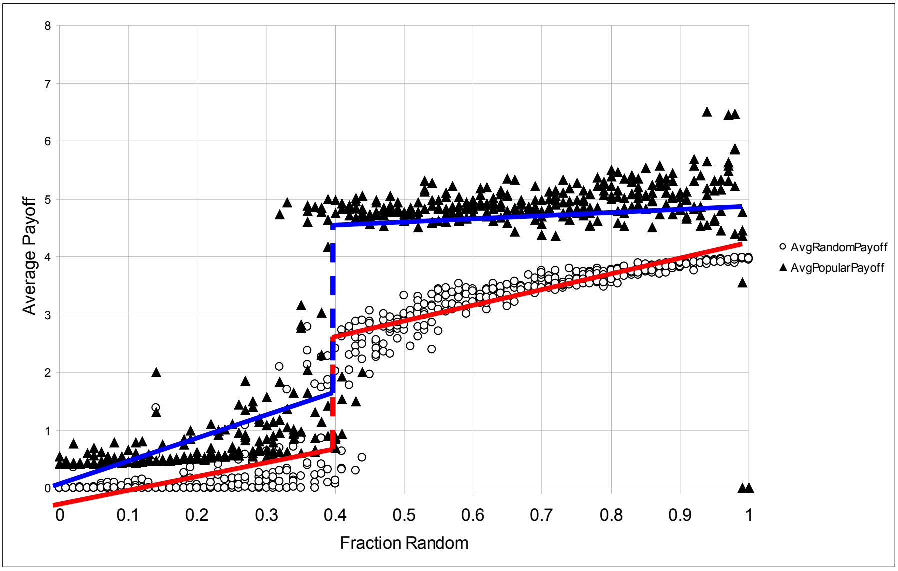

Graph 1: Random vs. Popularity Payoff Regressions

Random vs. Referral

The Random and Referral strategies were the second pair of heuristics I compared. Interestingly, this comparison yields similar results to the comparison between the Random and Popularity strategies. The Referral strategy yields higher payouts than the Random strategy, but the more agents that use Referral, the lower payouts are for all agents.

It should be noted that the variance in payoffs to Referral agents is significantly larger than the variance in the payoffs to Popularity agents in the previous comparison. This is clearly visible in Graphs 1 and 2, and is confirmed by the difference in the standard deviations of the average Popularity and Referral payoffs and the standard errors of the regressions. A possible explanation for this effect may be the "clustering" of agents using the Referral bias into groups with highly correlated payoffs, since groups will defect or cooperate together. Because grouped agents' payoffs will be correlated, the effective number of agents used in calculating average payoffs are lower, which raises variances.

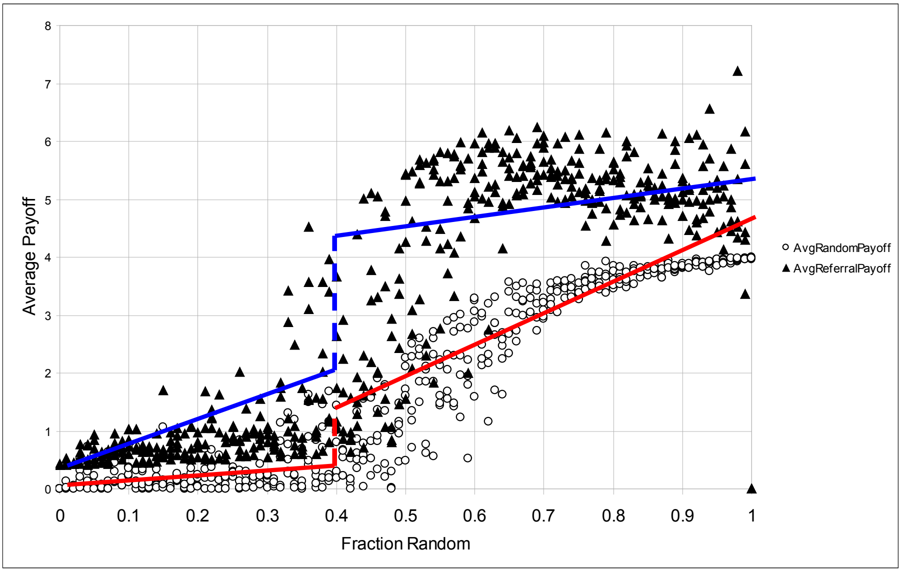

Graph 2: Random vs. Referral Payoff Regressions

Random vs. Unpopularity

There is relatively little difference between the Random and Unpopularity strategies in terms of their effect on either the number of cascades that occur, the length of those cascades, or the payouts realized by the agents. Cascades are relatively infrequent, occurring on the order of 1 or 2 times per 50,000 rounds, far fewer than the 8-10 cascades caused by the Popularity and Referral strategies over the same time period.

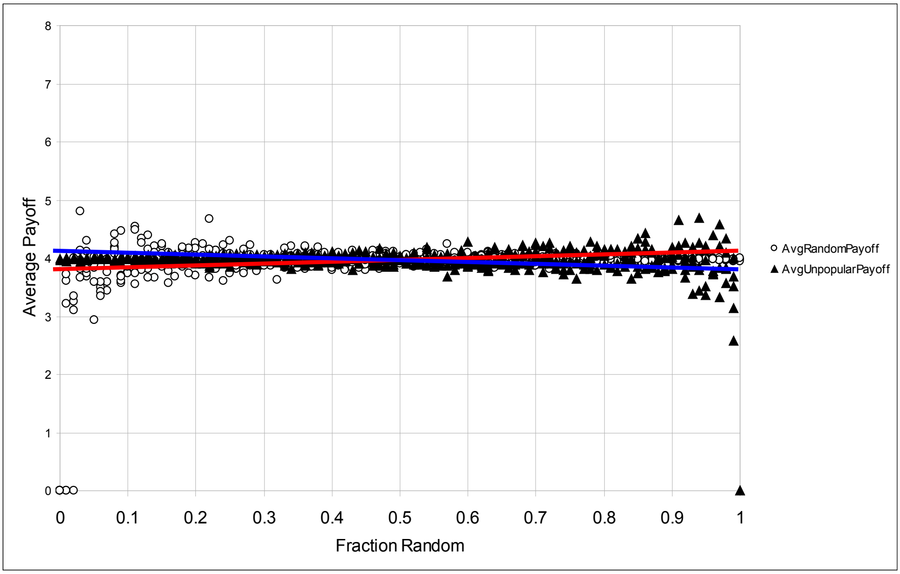

Graph 3: Random vs. Unpopularity Payoff Regressions

Random vs. Locality

The networks generated by agents playing the Locality strategy have the fewest cascades of any of the networks studied. Random dominates Locality as a strategy, most likely due to the ability of Random agents to "escape" local neighborhoods of defecting agents. The Random strategy also correlates with a small increase in the number of cascades, which decreases the average payout for both types when most agents play the Random strategy.

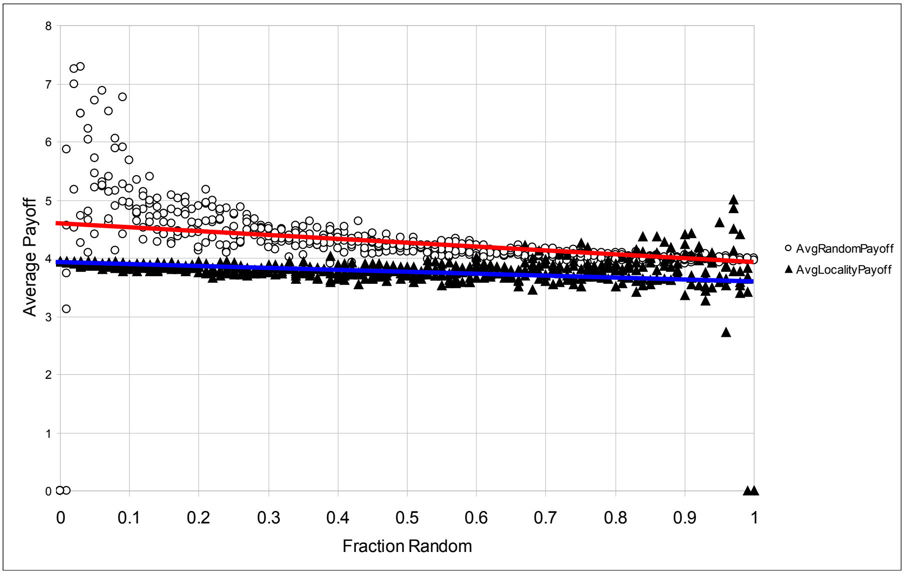

Graph 4: Random vs. Locality Payoff Regressions

Summary

The data indicate that some heuristics consistently outperform others, but may reduce the total payout of the network as a system when played by more than a small fraction of the agents.

On an individual level, the Popularity and Referral strategies outperform the Random strategy; the Unpopularity strategy produces the same outcomes as Random, and the Locality strategy performs worse than Random. The average payout for the entire network is maximized in the opposite order, however, with the Locality, Unpopularity and Random strategies benefiting the entire network more than the Popularity and Referral strategies, because Random, Unpopularity and Locality don't cause defection cascades.

V. Discussion

In this section, I give possible explanations for several of the phenomena observed in the data. I begin by discussing the similarity between the effects of the Referral and Popularity strategies, and move on to the individual benefits awarded by those strategies. I then discuss the global harms from those same strategies, and close by discussing the consequences of strategies which are individually beneficial but have negative externalities.

Similarity between Popularity and Referral

The similar performance of the Popularity and Referral strategies can be explained by the fact that agents are favored in proportion to their number of links in both cases. Even though the Referral strategy does not use the number of links as a direct criterion for selection, "popular" agents are more likely to be referred than other agents. As a result, while the Referral strategy increases the clustering of the network, it generally performs like the Popularity strategy, selecting agents with many neighbors more frequently than agents with few neighbors.

Why Do Popularity and Referral Increase Individual Payouts?

Despite the negative effects of the Popularity and Referral strategies, it is clear that agents playing these strategies receive higher payoffs than those playing Random. This is most likely because agents playing the Popularity and Referral strategies seem to accumulate more neighbors over time. Graph 5 is a histogram that shows the distribution of numbers of neighbors possessed by agents playing the Random and Popularity strategies in a representative run of the simulation with 70% Random agents and 30% Popularity agents. In this example, the Popularity agents average 4.8 neighbors each, while the Random agents average 3.7 neighbors. Because the number of neighbors agents have is directly related to their payouts (one point per neighbor in mutual cooperation), agents with more neighbors perform better.

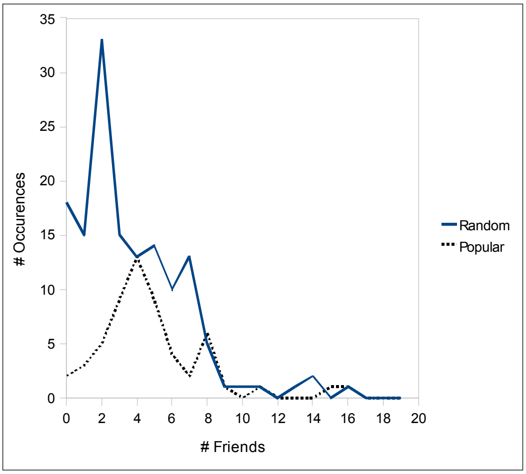

Chart 5: Random and Popularity Histograms for Number of Neighbors

How is it the case that agents who select partners based on Popularity themselves become more popular? This may be due to the correlation between playing Popularity or Referral and being less likely to defect when a neighbor defects. Figures 3 and 4 detail how this might work.

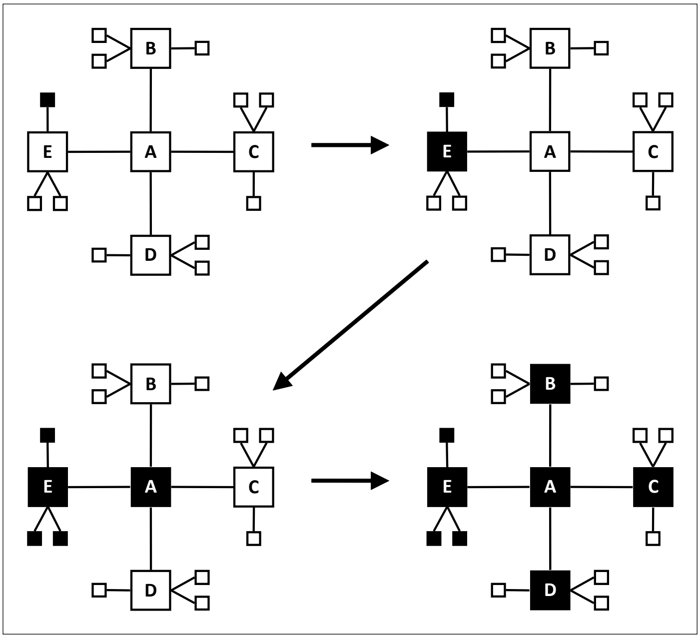

Figure 3: Infection Example without Role Model

Figure 3 is an example of how defection can spread through a network without a highly connected individual. After E defects in the example, A must choose whether or not to imitate E by defecting in the third round or imitate its other neighbors by continuing to cooperate. E receives a total payout of 4.5 (3 x T where T = 1.5). A's cooperating neighbors only get payouts of 4 (3 x R where R = 1) and A gets a payout of 3 (3 x R). As a result, A imitates E in the third round and defects.

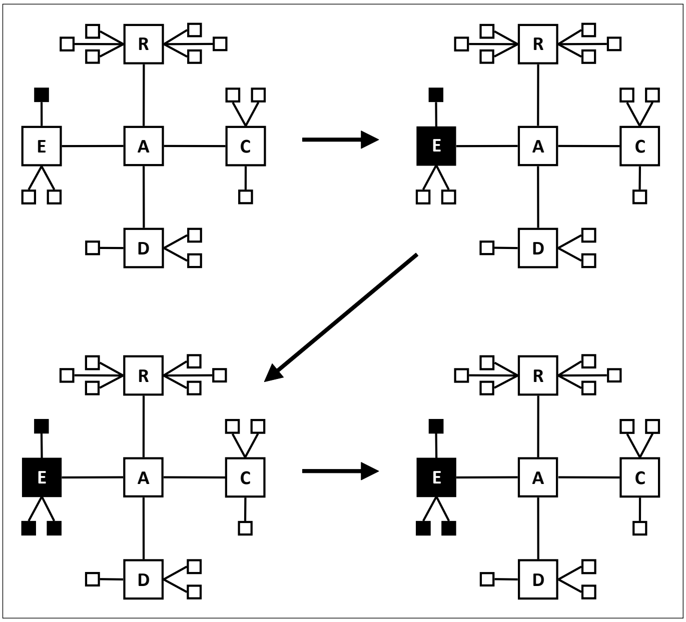

Figure 4: Infection Example with Role Model

Agents playing a Referral strategy are more likely to be connected to highly connected "role model" agents, represented by agent R in Figure 4. Role models reduce the likelihood of their neighbors converting when another neighbor defects. Because agents' total payoffs rise with the number of neighbors they have, R provides A with an example of a very high payoff (1 point from each of 7 neighbors, 7 total) to cooperation. E gets a higher payout from each of its parasitic relationships, but has fewer of them, yielding a total payout of 4.5 (1.5 points from each of 3 neighbors). A imitates R rather than E, because R's payout is greater than E's payout (7 > 4.5).

In the previous example without a role model, A will imitate any of its neighbors who defect, whereas when it has a role model, it will only imitate that role model. As a result, the frequency with which A defects is reduced.

Every time agents defect, their neighbors may choose to break links with them. As a result, agents who are able to defect less can accumulate more links over time, increasing their payouts. By connecting agents with role models, the Popularity and Referral strategies can improve the payouts of the agents that use them.

Rational, forward-looking agents would want to accumulate as many links as possible to increase their payoffs, and would therefore resist the temptation to defect in order to encourage their neighbors not to abandon them. While agents playing Popularity and Referral don't work through this logic explicitly, their imitation and friend choice heuristics cause them to take actions that appear to be fully rational.

How Do Popularity and Referral Harm Global Payouts?

While Popularity and Referral are each beneficial in their own respective manner, they have different effects on the network as a whole. The more agents that use the Popularity and Referral strategies, the worse the average payout for the entire population. This is due to the increasing prevalence of defection cascades as the proportion of agents using Popularity and Referral increases.

This may be caused in large part by the exclusion of a large number of agents from the network as a side effect of the Popularity and Referral strategies. When these strategies are played, many simulated agents gradually lose all of their links and no longer interact with the social network. Because agents using the Popularity and Referral strategies favor other agents with many neighbors, agents with few neighbors may lose neighbors faster than they regain them, until they are completely cut off from the network. Because the Random, Unpopularity and Locality strategies do not favor agents with many neighbors, these strategies do not have this ostracizing effect. The proportion of agents that are ostracized increases as more agents use Popularity and Referral, a relationship confirmed by Chart 6. In particular, the proportion of agents ostracized spikes when more than 60% of agents play the Random or Referral strategy, which coincides with the structural break in the payoff and cascade data.

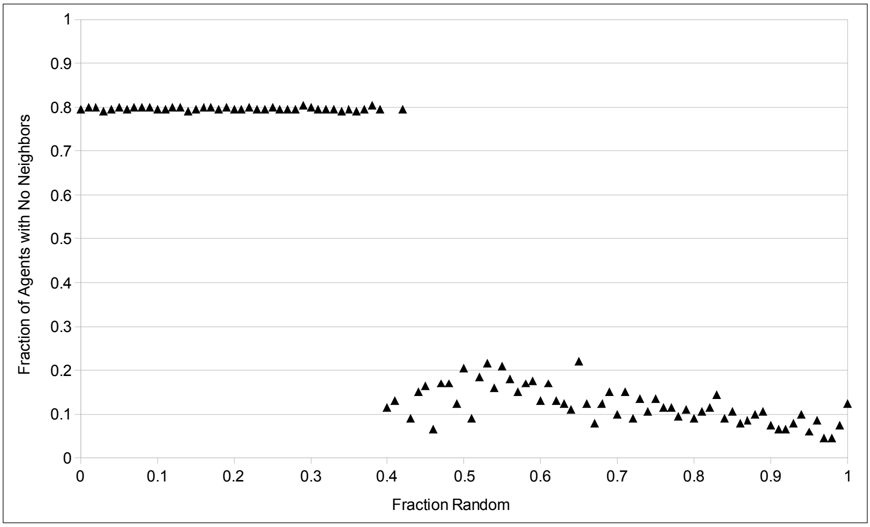

Chart 6: Random vs. Popular: Number of Agents without Neighbors

One of the properties of the model is a conservation of the number of links. This means that as agents without neighbors accumulate, a shrinking pool of agents uses all of the links in the network. While the average number of links possessed by each agent in the simulation is constant due to the conservation of links, the average number of links possessed by each agent connected to the network increases as more agents are ejected from it.

As a result, networks such as these are more susceptible to defection cascades. Defection cascades spread faster and more completely because of the decreased network diameter and increased rate of initial exponential growth when agents have more links. This explains the upward trend in cascade rates as more agents use the Popularity and Referral biases.

When more than a certain fraction of agents use the Popularity or Referral strategies, the system enters a state of nearly continuous defection. The critical fractions of agents seem to be about 60% Popularity strategy or about 55% Referral strategy. When too many agents play the Popularity or Referral strategies, the network is unable to recover from defection cascades, and continually reshuffles the links as defecting agents end relationships with other defecting agents. As a result, the network shrinks down to a highly connected cluster of 40-50 agents, and the majority of agents are excluded from the network with no neighbors. The social network cannot recover from this disintegration, an event comparative to phase changes often observed in physics.

Prisoner's Dilemma within a Prisoner's Dilemma

It is clear from the results that the Popularity and Referral strategies dominate the Random and Unpopularity strategies, which in turn dominate the Locality strategy. In terms of social optima, however, the strategies are reversed. The Locality strategy is best for the network as a whole, and the Popularity and Referral strategies worst, due to their respective effects on the frequency and severity of defection cascades.

This means that agents choosing a friendship choice heuristic face a Prisoner's Dilemma of sorts. They can choose the Popularity or Referral strategies, benefiting themselves but harming the network, or they can choose Random, Unpopularity or Locality, which neither benefit the agent nor harm the network. The Social Optimum occurs when all agents choose Random, Unpopularity or Local, but the Nash Equilibrium occurs when all agents choose Popularity or Referral. To the detriment of the entire network, utility maximizing agents left to their own devices will choose to use the Popularity and Referral strategies to select their friends.

VI. Conclusion

I have shown that two link creation heuristics, the Popularity and Referral strategies, have interesting emergent properties when played by a significant proportion of the agents in a dynamic local interaction network. New partners selected using a Popularity or Referral strategy are more likely to be cooperative; therefore, agents that use these strategies consistently have higher payoffs than agents that play the Random strategy. At the same time, the more agents that play the Popularity and Referral strategies, the more frequently defection cascades occur, which lowers the average payouts of every agent in the network.

This emergent Prisoner's Dilemma is significant because it is not obvious when studying the friendship choice heuristics analytically. Real life examples of cascades like financial contagion are rare and therefore difficult to study empirically. Backward-looking models do a poor job of predicting cascade behavior due to the lack of appropriate data necessary to calibrate the model. Dependence on these models has been the downfall of many investment companies, most notoriously Long Term Capital Management. It is worth considering whether our financial institutions create incentives for firms to use strategies that have negative externalities that destabilize networks. Hopefully someday we will have sufficient understanding of financial contagion to prevent this.

VII. Acknowledgements

Many thanks to my thesis advisor, Professor Nathan Larson, and the director of the Distinguished Majors Program, Professor William Johnson. This project would not have been possible without their help and guidance. I am also grateful to the University of Virginia Computer Science Department for its generosity in allowing me to use the Centurion cluster to run my simulation. Finally, thank you to my wife Michelle and to the editors of the Michigan Journal of Business for their substantial aid in revising this document for publication.

References

Axelrod, Robert. The Evolution of Cooperation. New York: Basic Books, 1984.

Axelrod, Robert, and Dion, Douglas. "The Further Evolution of Cooperation." Science 242 no. 4884 (1988):1385.

Axelrod, Robert, and Hamilton, William Donald. The Evolution of Cooperation. Science 211, no. 4489 (1981): 1390.

Axtell, Robert, Epstein, Joshua, and H. Peyton Young. "The Emergence of Classes in a Multi-Agent Bargaining Model," in Generative Social Science, edited by Joshua Epstein. Princeton: Princeton University Press, 2006.

Becker, Gary. "Irrational Behavior and Economic Theory." The Journal of Political Economy 70 no. 1 (1962): 1-13.

Beinhocker, Eric. The Origin of Wealth. Boston: Harvard Business School Press, 2006.

Bergstrom, Theodore. "Evolution of Social Behavior: Individual and Group Selection." Journal of Economic Perspectives 16 no.2 (2002):67-88.

Buchanan, Mark. Nexus. New York: Norton & Co, 2002.

Dodds, Peter, Muhamad, Roby, Watts, Duncan. "An Experimental Study of Search in Global Social Networks." Science 301 no. 5634 (2003): 827-9.

Eshel, Ilan, Samuelson, Larry, and Shaked, Avner. "Altruists, Egoists, and Hooligans in a Local Interaction Model." The American Economic Review 88 no. 1 (1998).

Farmer, J. Doyne. "Toward Agent-Based Models for Investment." Benchmarks and Attribution Analysis. The Association for Investment Management and Research (2001): 61-70.

Gladwell, Malcolm. The Tipping Point. New York: Little, Brown and Co, 2000.

Hanaki, N., Peterhansl, A., Dodds, P., Watts, D. "Cooperation in Evolving Social Networks." Management Science 53 no. 7 (2007):1036-50

Levine, David, and Pesendorfer, Wolfgang. "The Evolution of Cooperation Through Imitation." Games and Economic Behavior 58 no. 2 (2007): 293-315.

Johnson, Steven. Emergence. New York: Scribner, 2001.

Larson, N. "Match Choice and Ghettoization in evolutionary games." Journal of Economic Theory 117 (2004):1-28.

Lee, In Ho, and Valentinyi, Akos. "Noisy Contagion Without Mutation." Review of Economic Studies 67 (2007):47-56.

Lewin, Roger. Complexity: Life at the Edge of Chaos. Chicago: University of Chicago Press, 2000.

Lindgren, Kristian, and Nordahl, Mats. “Cooperation and Community Structure in Artificial Ecosystems,” in Artificial Life: An Overview, edited by Christopher Langton, 15-38. Cambridge: MIT Press, 1995.

Milgram, Stanley. “The Lost Letter Technique: A Tool of Social Research.” The Public Opinion Quarterly 29 no. 3 (1965): 437-8.

Milgram, Stanley. “The Small World Problem.” Psychology Today 2 (1967): 60-67.

Morris, Stephen. “Contagion.” Review of Economic Studies 67 (2000): 57-78.

Nelson, Richard, and Winter, Sidney. An Evolutionary Theory of Economic Change. Boston: Belknap Press of Harvard University, 1982

Schelling, Thomas. “Models of Segregation.” American Economic Review 68, no. 3 (1971): 488-93.

Watts, Duncan. “The ‘New’ Science of Networks.” Annual Review of Sociology 30 (2004):243-70.

Watts, Duncan, Dodds, Peter, Newman, M. E. J. “Identity and Search in Social Networks.” Science 296 no. 5571 (2002): 1302-5.

Watts, D. “A simple model of global cascades on random networks.” Proceedings of the National Academy of Sciences of the United States of America 99 no.9 (2002): 5766-71.

Watts, D., Strogatz, S. Collective dynamics of ‘small-world’ networks. Nature 393 no. 684 (1998): 409-10.

Wynne-Edwards, Vero Copner. Animal Dispersion in Relation to Social Behavior. London: Oliver & Boyd, 1962.