Phase transitions in social networks

Jahan Claes

Abstract

In both evolution and economics, populations sometimes cooperate in ways that do not benefit the individual, and sometimes fail to cooperate in ways that would benefit the population. One method to explain the success or failure of cooperation is the social network model, in which simplified two-person interactions take place over a network of social connections. In this paper, I explore some of the experimental evidence of cooperation in microorganisms and humans. Then, I explain the theoretical models used to predict when cooperation will occur, with particular attention to the prediction of phase transitions.

1 Introduction

In economics and evolutionary theory, the existence of cooperative behavior in large groups has been a longstanding puzzle. Cooperation, while mutually beneficial to society, frequently comes with a temptation for any individual to defect. How, then, does large scale cooperation exist? Why does the United States have public radio and privately run homeless shelters? Why do some animals, such as guppies, defend each other against predators when it would be safer to run away? And why does cooperation sometimes break down? Why does the United States not have free grocery stores, or a communal pool of cars that anyone can use? Why do other animals leave their friends to be eaten?

One method to explain large scale cooperation is the social network model, in which a population interacts with some fixed set of partners via a prisoner's dilemma type game. In these models, both the rules govering the interaction between partners and the overall structure of society affects the presence or absence of cooperation. As we will see, these models can make powerful predictions, but are sensitive to the details of the network they are built on.

In this paper, we will focus on the presence of phase transitions in these social networks, as one tunes the parameters governing the interactions in our model society. We will see that the cooperative behavior of a social network can display discontinuous and continuous phase transitions, as well as a buildup of critical fluctuations.

The overall outline of the paper is as follows. First, we will describe some experimental studies of cooperative behavior. Next, we will describe in detail the social network model of cooperative behavior. Finally, we will discuss the theoretical predictions of phase transitions in these models, both in mean field theory and beyond mean field theory.

2 Experimental results

Here, we describe some experimental results on prisoners dilemma type problems. This discussion is by no means comprehensive, as a great deal of research has been done on the prisoner's dilemma, and a complete overview of the literature would be beside the point. Instead, we focus on a few examples, which serve to illustrate the importance of both the interactions between individuals and the overall social structure.

In economics, [1] demonstrate human cooperation via altruistic punishment. In that experiment, players were allowed to put money in to a public pool, where it would be doubled and then divided up evenly among all players. In this game, it was found that participants would pay money to punish defectors, even if the punishment did not benefit the participant. Thus, most humans were willing to behave cooperatively even if it did not directly benefit them and would not benefit them in the future.

In evolutionary biology, one might expect organisms to evolve to benefit the species as a whole; however, if a random mutation allows a single organism to take advantage of the species' cooperation, that organism will likely breed faster and dominate the species. This was demonstrated in [2], in which it was shown that for high enough densities of a certain RNA virus, the virus would eventually evolve towards non-cooperation. However, this paper also showed that virus strains which evolved in low-density environments remained cooperative, because subsequent generations were more genetically similar to the ancestors. This is a first suggestion that the geometry of the social network is important: high density allows for many interactions and a dense social network, while low density creates a sparse social network.

Similarly, [3] demonstrates the evolution of non-cooperation in yeast. However, in this study, it was found that unlike in the RNA virus, the non-cooperator did not dominate the population; instead, the cooperators and non-cooperators coexisted at some fixed fraction. The authors show that this is due to the fact that the yeast is not in a perfect prisoner's dilemma situation, in which defecting is always advantageous. Rather, after a certain fraction of the population has defected, the cost to defect is higher than the cost of cooperation. This variation of the prisoner's dilemma, the so-called **snowdrift** game, will be discussed in more detail in section 3.1 below. For now, the important takeaway is that the rules of interaction affected the macroscopic fraction of cooperators.

We thus see that the cooperative outcome of the social game can be expected to depend on both the pairwise interactions and the overall social structure.

3 The social network model of cooperative games

In this section, we explain the model used to describe cooperative games, the social network model, first described in [4]. The basic framework the social network model is that the population is made up of N individuals, who plays a collection of two person games with some fixed set of partners in the population whom they are connected to. After all games are completed, the successful individuals either reproduce or convince others to adopt their strategy, while the unsuccessful individuals die off or are convinced to adopt a more successful strategy. The game then iterates.

3.1 The two person cooperative game

Our two person game will be a prisoner's dilemma type game, where each player has only two strategies available to them. The may defect (D) or cooperate (C). The player receives a payoff depending on both his strategy and his opponents strategy. The table below defines the outcome of the game. For each strategy of the player on the right and opponent's strategy above, the player receives a the payoff given in the table. For example, if a player chooses (C) while the opponent chooses (D), the player receives a payoff of c. Note that we can always rescale our units so that the difference between the (C,C) payoff and the (D,D) payoff is 1. In addition, all the results we describe below depend only on the difference between payoffs, so we may set the (D,D) payoff to 0 without affecting outcomes. Thus, the table given below is completely general. The matrix describing the payoff table is frequently denoted by $\mathbf{P}$.

		Opponent	
		C	D
Player	C	1	c
	D	b	0

$$\mathbf{P} = \begin{pmatrix} 1 & c \\ b & 0 \end{pmatrix}$$

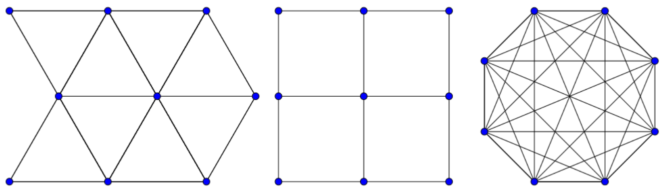

Figure 1: Three examples of social networks. From left to right: a triangular lattice, a square lattice, and a complete graph.

We can classify the games into four main categories, depending on the values of b and c [5, 6]. The case where $b > 1$, $c < 0$, is the classic **prisoner's dilemma** game (PD). No matter the strategy used by the opponent, it is optimal for the player to choose to defect (D). Thus, one expects every player to defect and for every player to receive a payoff of 0, despite the fact that mutual cooperation could lead to everyone receiving a payoff of 1.

The case where $b > 1$, $c > 0$, is called a **snowdrift** game (SD). In this game, like the PD, each player wants to be the only one to defect. However, unlike the PD, if a player expects their opponent to defect, it is rational for them to choose (C). Thus, one expects players to attempt to choose the opposite strategy as their opponent.

The case where $b < 1$, $c > 0$ is called a **harmony game** (HG). In this game, defecting always gives a worse outcome, thus one expects players to always choose (D).

Finally, the case where $b < 1$, $c < 0$ is called a **stag hunt game** (SH). In this game, one has no incentive to defect unless one believes their opponent will defect. Thus, one expects players to choose (C,C), but to choose (D,D) if they do not trust their opponent to be rational.

3.2 Cooperative games and geometry

To describe population-wide dynamics, we need to describe both the two-player game being played, and the pairs of people playing it. Given a population of size N, we'll describe the pairs of players by a graph on N vertices. Two people in the population play a round of our game if they are connected by a node on the graph. These graphs are called **social networks**, as they are meant to encode which people interact in our model society [5].

Figure 3.2 gives three examples of possible social networks. In the triangular lattice, each person interacts with their six nearest neighbors, while in the square lattice each person interacts with their four nearest neighbors. In the complete graph, each person interacts with every other member of the population.

We should take special notice of the complete graph. Since each person is connected to every other person, each person receives a payoff which depends only on their strategy and the average of all other people's strategies. Thus, for large N, this social network corresponds to **mean field theory** on a given social game [5].

3.3 Time dynamics of cooperative games

The final ingredient we need for our theoretical description is a rule for how people's strategies change with time. Intuitively, we want to demand that after each round of the cooperative game, successful strategies proliferate while unsuccessful strategies die out. We might also wish to introduce some noise, to represent random mutations in strategy. Here we introduce a few possible update rules that are used below; obviously, others are possible.

Denote the strategy of player i by a vector $\vec{s}_i$, where $\vec{s}_i = \binom{1}{0}$ if player i is cooperating, and $\vec{s}_i = \binom{0}{1}$ otherwise. Then after one round of the game, player i receives a total payoff of

$$P_i = \sum_j \vec{s}_j^{\,T} \mathbf{P} \vec{s}_i$$

where the sum is over the neighbors j of i in the social network.

Our simplest update rule is [7]: i chooses a neighbor j at random. If $P_i > P_j$, i does nothing, since his strategy is superior. If $P_i < P_j$, P_i adopts strategy $\vec{s}_j$ with probability

$$W[\vec{s}_i \to \vec{s}_j] = \eta(P_j - P_i) \qquad (1)$$

where η is some constant that ensures $\eta(P_j - P_i) < 1$ for all possible P_i, P_j. The constant η affects the rate at which strategies can spread through a social network, but generally does not affect the long-time behavior.

In the case of a complete graph with $N >> 1$ players, the only relevant quantities are the fractions p_C and p_D of cooperators and defectors. In this mean field case, equation 1 reduces to [6, 7]

$$\frac{dp_C(t)}{dt} = p_C(t) \left[\sum_{\beta=C,D} \mathbf{P}_{C\beta} p_\beta(t) - \sum_{\alpha,\beta=C,D} p_\alpha(t) \mathbf{P}_{\alpha\beta} p_\beta(t) \right] \qquad (2)$$

and of course $p_D = (1 - p_C)$. This equation is called the **replicator equation**. It says that the rate of change of strategy C is proportional to the current fraction of people using C times the difference between payoff for C and the average payoff in the population. If C has a greater average payoff than the population its use will increase; otherwise its use will decrease. Note that η from equation 1 has disappeared, as we've absorbed it into our parameter t.

Finally, we will use a modified update rule that includes random mutation. In the update rule given by equation 1, we specified that a player would never change to a strategy that was worse than their current strategy. In this update rule, we introduce a noise parameter that allows a player to occasionally adopt a worse strategy. Each player i randomly chooses a neighbor j, and switches to strategy $\vec{s}_j$ with probability [8]

$$W[\vec{s}_i \to \vec{s}_j] = \left[1 + e^{\frac{P_i - P_j}{K}} \right]^{-1} \qquad (3)$$

Here K is a parameter that determines the noise in the system. If $K = 0$, the neighboring strategy is adopted whenever $P_j > P_i$, while $K > 0$ there is some probability of i declining to switch even when $P_j > P_i$ and some probability of switching even when $P_j < P_i$.

The mean field theory version for of this equation [8]

$$\frac{dp_C(t)}{dt} = p_C(t)[1 - p_C(t)] \tanh\left(\frac{[b + c - 1]p_C(t) - c}{2K}\right) \tag{4}$$

although we will see that this equation predicts the same long-time behavior as equation 2.

4 Phase transitions in social networks

Here, we discuss the phase transitions predicted in social networks. While a complete theory on general graphs is not known, the phase diagram is essentially complete at a mean field level. Beyond mean field, we will outline the results that have been studied on specific networks.

4.1 Phase transitions in mean field theory

We will explore the mean field theory given by the replicator equation 2. To solve for the long time behavior, we want to find the stationary solutions and determine their stability. The stable solutions correspond to the long-time behavior of the game. We will characterize the phase of the system by the values of the stationary solutions of $p_C \equiv p$. The equation for p can be written as

$$0 = p(1 - p)[c(1 - p) - (b - 1)p] \tag{5}$$

The stationary solutions are then given by $p_1 = 0$, $p_2 = 1$, and $p_3 = \frac{c}{c+b-1}$. Note that the stability of these solutions depends on the values of c and b, and that the solution p_3 is only a valid solution when $0 \leq p_3 \leq 1$.

We can classify the phases of our game using the nomenclature of section 3.1 based on the stable and unstable solutions [6]

	$c > 0$	$c < 0$
b>1	**SD** p_1 unstable p_2 unstable p_3 stable	**PD** p_1 stable p_2 unstable p_3 invalid
b<1	**HG** p_1 unstable p_2 stable p_3 invalid	**SH** p_1 stable p_2 stable p_3 unstable

The case of the stag hunt (SH) game is particularly interesting: here, there are two stable solutions. The system evolves towards total cooperation if it initially satisfies $p_c(0) > p_3$, and evolves towards total defection otherwise.

We see from this classification that the transitions between certain games must be discontinuous. For example, if we tune b and c to go directly from PD to HG, the order parameter must jump from $p = p_1$ to $p = p_2$. Other transitions may be continuous; for example, the transition from PD to SD is continuous, since at the phase boundary $c = 0$, $p_3 = p_1$.

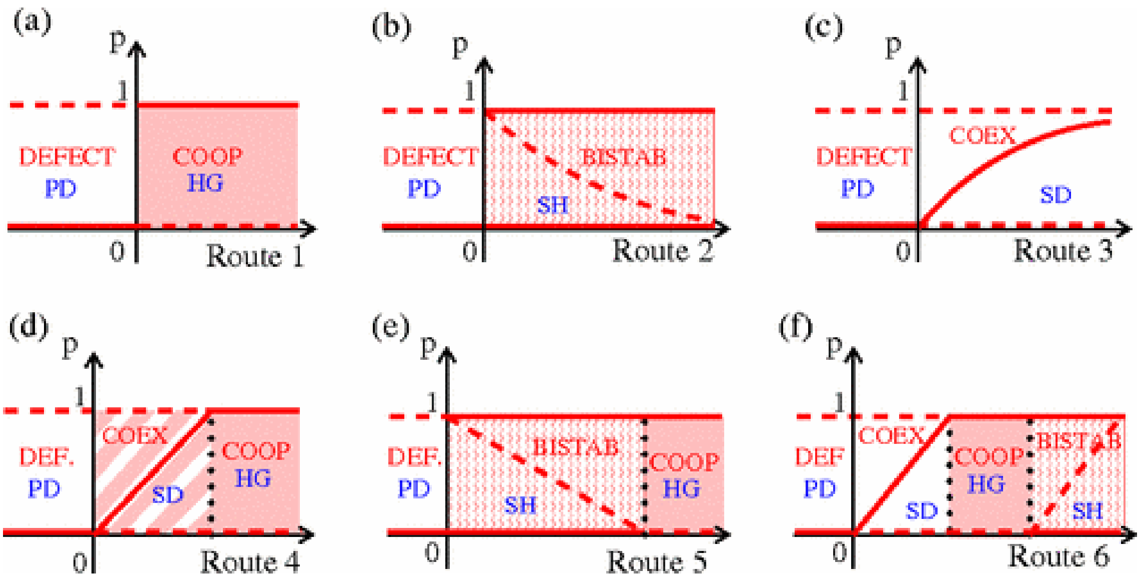

Figure 2: Figure from [6]. Each graph describes a certain path in phase space. Solid lines represent stable equilibrium, while dotted lines represent unstable equilibrium. The white regions of the graphs are the PD phase, where defection occurs. The solid red regions are the HG phase, where cooperation occurs. The striped regions are the SD phase, where cooperation and defection coexist. Finally, the dotted red regions are the SH phase, where both cooperation and defection are stable, and are separated by an unstable coexistence line.

In [6], the authors illustrate the possible phase transitions from a PD to a game in which $p \neq 0$. Route 1 illustrates the transition between PD and HG, in which p changes discontinuously. Route 2 illustrates a transition between PD and SH, where $p = 0$ the entire time unless the parameters are tuned to make $p_3 << 1$, in which case small fluctuations can drive the system above p_3 and thus into the $p = 1$ phase. Route 3 demonstrates the continuous phase transition between PD and SD. Route 4 illustrates a route to a HG phase than is a continuous transition rather than a first-order transition, where we first tune through a SD phase. Route 5 illustrates the transition PD-SH-HG. In this transition, it is not clear where p jumps from $p = 0$ to $p = 1$; the jump must happen wherever the amplitude of small fluctuations exceeds p_3 and drives the system to p_2. However, we know by the time the system reaches HG the system must have $p = 1$. Finally, route 6 illustrates the continuous transition PD-SD-HG-SH. This transition guarantees that the SH game ends up in the $p = 1$ cooperative phase by first going through a HG phase.

Obviously, one could consider many other paths between phases, but all the paths can be characterized entirely by the stability of p_1, p_2, and p_3 and the value of p_3. Note that even at the mean field level we see nontrivial phase boundaries. These phase boundaries have nontrivial implications. For example, if a government wanted to transition a certain segment of society between the PD and SH phase, it could pick many possible paths. However, if the government wants to avoid economic shocks it should avoid discontinuous transitions, and if it wants to ensure cooperation it should end in the $p = 1$ equilibrium of SH. Thus, the government should follow route 6 rather than route 2.

We can also consider playing the same game with equation 4. While this equation

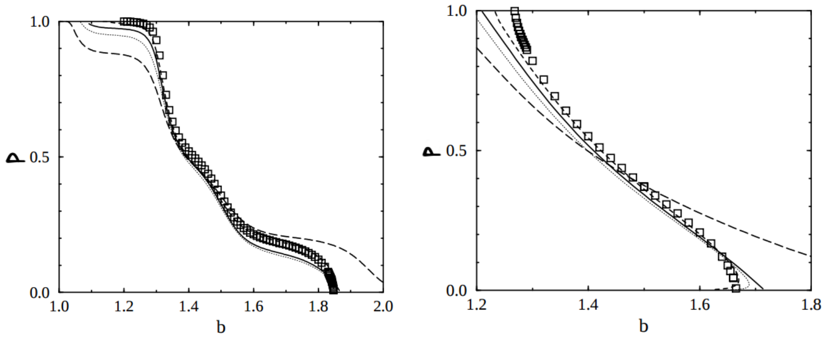

Figure 3: Figures from [9]. p as a function of b on the square lattice. Left: $K = .1$. Right: $K = .5$. In both cases, we see that the presence of the lattice has allowed cooperation to persist into the PD regime. The squares are the Monte Carlo data, while the dotted lines represent a more sophisticated mean field approximations generated by clustering groups of nodes in the lattice.

has an additional parameter K to tune, this parameter does not change the mean field equilibrium solutions. It is straightforward to check that the solutions p_1, p_2, and p_3 still apply, and that the stability is identical as well. Thus, we expect the same phase diagram for this equation.

4.2 Beyond mean field theory

Here, we review the work that has been done on specific lattices beyond mean field theory. Broadly, we find that phase boundaries are modified towards more cooperation on these lattices vs the mean field theory. Intuitively, this is because defectors cannot spread arbitrarily; instead, they quickly surround themselves with other defectors and reduce their payoff compared to the cooperators. This phenomena, in which the presence of a lattice enhances cooperation, is called **lattice reciprocity**.

In the following sections, we show that Monte Carlo results on a simple lattice modifies the phase boundary from mean field theory, while mean field theory neglects a phase boundary entirely in a more exotic social network due to critical fluctuations. Finally, we explore the effect of noise on the phase boundaries, which cannot be captured by the mean field theory.

4.2.1 Mean field theory vs Monte Carlo on a square lattice

The simplest lattice we may explore is the square lattice. Here, we consider the limit of c small and negative, and look at the behavior of p as a function of b. We'll use the update given in 3, so that a parameter K describes the noise in the system.

Because $c < 0$, we're tuning between SH and PD. The mean field theory predicts either a continuous transition with $p = 0$ or a discontinuous transition from $p = 1$ to $p = 0$ when tuning across the phase boundary $b = 1$. In [9], the authors used Monte Carlo methods to determine the p as a function of b. Their results are plotted in figure 4.2.1. They found that the presence of the lattice increased the critical value of b and made the phase transition continuous, while noting that tuning K towards zero seemed to increase the sharpness of the transition.

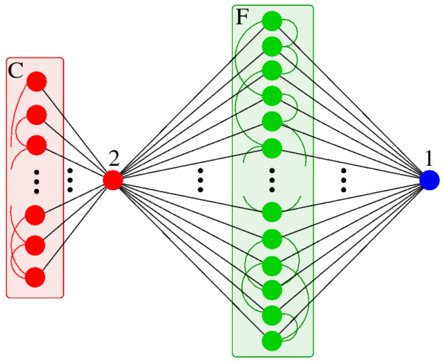

Figure 4: Figure from [7]. The big brothers social network. The big brothers, labeled as 1 and 2, are connected to groups C and F. C and F are regular graphs, in that each node is connected to k other nodes.

We have thus seen our first example of lattice reciprocity, as the lattice allowed cooperation to persist in the PD regime.

4.2.2 Critical fluctuations in a big brother network

Here, we present results from [7] which shows a network that has a phase transition mean field theory fails to capture. In this model, we use the update given by equation 1, and study an exotic network, called the **dipole** or **big brothers** network, shown in figure 4. Here, two "big brothers", labeled 1 and 2, are connected to groups of nodes C and F. Each of these groups is a regular graph, in that every node is connected to k other nodes for some fixed k. We denote by n_c and n_F the number of nodes in C and F. This graph represents a society in which there exist two main dominant figures, and two distinct social groups.

In this network, the authors performed a mean field calculation in which fluctuation in the groups C and F were neglected, but the evolution of nodes 1 and 2 were treated exactly. They then compared this result to Monte Carlo simulations for $n_F = 4000$ and decreasing values of n_C. The results are plotted in the first graph of figure 4.2.2. They find that the Monte Carlo results show that p goes linearly to 0 in the limit $n_C/n_F \to 0$, while the mean field calculations show p asymptotically decaying as $b \to \infty$. Thus the mean field theory has entirely missed a phase transition. This cannot be thought of purely as a failure of mean field theory for small n_C, since it is possible to take the $n_C/n_F \to 0$ limit without n_C itself being small.

The mean field theory itself predicts the fluctuations to smoothly decrease in this critical region. To explain this breakdown of mean field theory, the authors also computed the fluctuations $\langle\sigma\rangle = \langle p^2\rangle - \langle p\rangle^2$ in their Monte Carlo simulation. These results are shown in the second graph of figure 4.2.2. We see that as $n_C/n_F \to 0$, the fluctuations diverge near the critical value of b. We thus see that this breakdown of mean field theory is directly related to the buildup of fluctuations near the critical point.

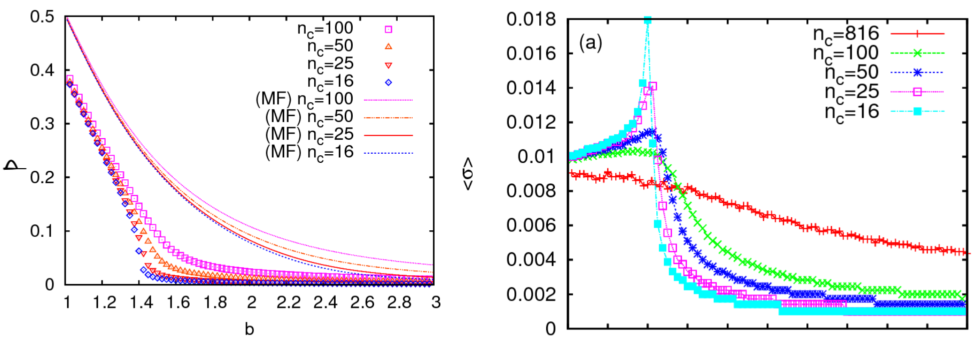

Figure 5: Figures from [7]. Left: p as a function of b in Monte Carlo and mean field calculations. Monte Carlo predicts a sharp transition to $p = 0$ when $n_C/n_F \approx 0$, while mean field theory does not. Right: Buildup of critical fluctuations as $n_C/n_F \to 0$.

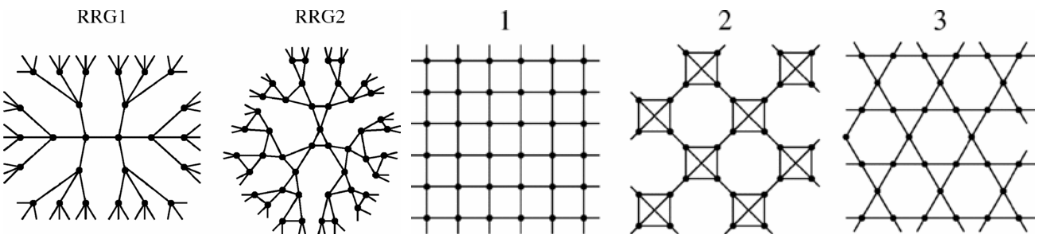

Figure 6: Figures from [8, 10]. The five lattices considered in section 4.2.3. From left to right: two random regular graphs, a square lattice, a four-site clique lattice, and a kagome lattice

4.2.3 Social networks and noise

We want to characterize the dependence of cooperation on noise. Does noise/randomness increase cooperation or decrease it? As we've seen, mean field theory predicts the noise parameter K has no effect on the presence or absence of cooperation. Here, we explore the effect of noise on five different model lattices, shown in figure 6. Since we are working with a noise parameter, our update is given by equation 3.

We restrict to small but negative c, so we are considering the PD and SH games. The fact that c is small implies that fluctuations should send us to the $p = 1$ phase of the SH game, at least according to mean field theory. Generally, then, we expect that for fixed K, as we tune from $b << 1$ to $b >> 1$ we will see a transition from $p \neq 0$ to $p = 0$ at some critical value b_{cr}. Mean field theory predicts $b_{cr} = 1$. In [8] and [10], the authors used Monte Carlo simulations to determine the actual critical values of b. The results are plotted in figure 7. We see that the effect of noise is to encourage cooperation compared to $K = 0$ on the first random regular graph, square lattice, and clique lattice, but discourage cooperation on the second random regular graph and the kagome lattice. Thus, noise does not have a predictable effect on cooperation; it depends on the social network studied. We also note that for all of the lattices, the presence of the lattice increased the critical value compared to the mean field prediction of $b_{cr} = 1$, so all of these lattices exhibit lattice reciprocity.

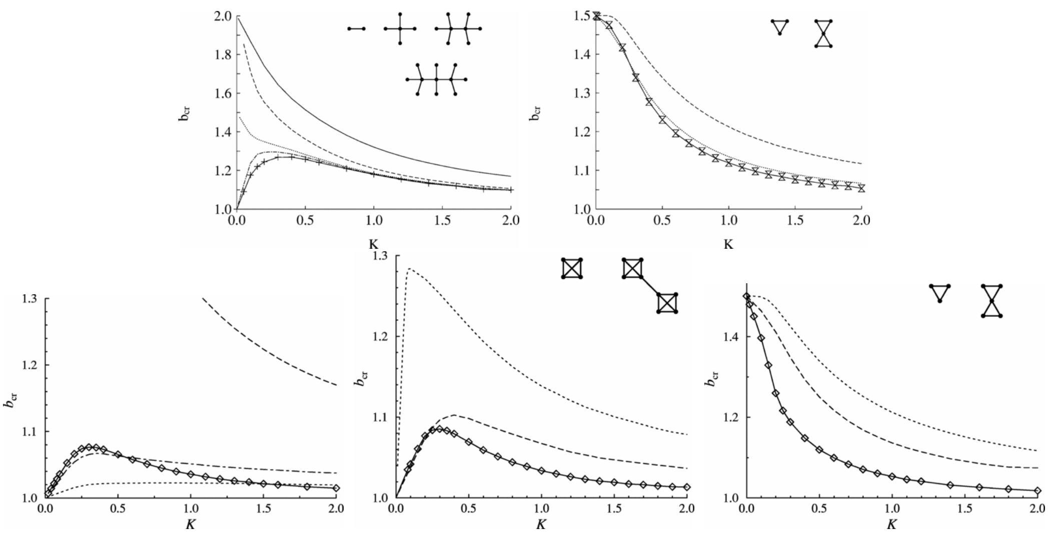

Figure 7: Figures from [8, 10]. From top left to bottom right: the critical values b_{cr} for lattices RRG1, RRG2, 1, 2 and 3 of figure 6. The data points connected by solid lines are the Monte Carlo calculations, while the dotted lines are again more sophisticated mean field calculations.

5 Conclusion

There are, of course, many other studies of social networks that were not mentioned in this paper, as they did not deal with the phase transition behavior. An excellent and extensive review article from a physics perspective is [5].

Social network models demonstrate that the presence or absence of cooperative behavior depends on both the type of payoffs considered and the overall structure of society. Moreover, the paths between cooperation and defection can cross nontrivial phase boundaries, in which singular behavior and critical fluctuations are observed. While these phases and phase boundaries are understood at a mean field level, the mean field theory can break down on more complex social networks. So far, there is no general theory connecting the structure of these social networks to the phases of social cooperation on them, or a general theory of what parameters might enhance social cooperation (besides the obvious parameters in the payoff matrix). One striking example of this is the effect of K on the phase boundary between PD and SH. Depending on the social network, we saw that K can either increase or decrease cooperation, and there is no obvious way to predict the effect from the geometry of the network. If the microscopic structure of these networks significantly affects their predictions, this limits their applicability to predict experiment, since the underlying social network is not known (and may not even rigidly exist). Thus, further work is needed to extract meaningful results that do not depend on the microscopic details of the network.

References

[1] Ernst Fehr and Simon Gächter. Altruistic punishment in humans. *Nature*, 415(6868):137–140, 2002.

[2] Paul E Turner and Lin Chao. Prisoner's dilemma in an rna virus. *Nature*, 398(6726):441–443, 1999.

[3] Jeff Gore, Hyun Youk, and Alexander Van Oudenaarden. Snowdrift game dynamics and facultative cheating in yeast. *Nature*, 459(7244):253–256, 2009.

[4] Martin A Nowak and Robert M May. Evolutionary games and spatial chaos. *Nature*, 359(6398):826–829, 1992.

[5] György Szabó and Gabor Fath. Evolutionary games on graphs. *Physics reports*, 446(4):97–216, 2007.

[6] Dirk Helbing and Sergi Lozano. Phase transitions to cooperation in the prisoner's dilemma. *Physical Review E*, 81(5):057102, 2010.

[7] LM Floría, C Gracia-Lázaro, J Gómez-Gardenes, and Y Moreno. Social network reciprocity as a phase transition in evolutionary cooperation. *Physical Review E*, 79(2):026106, 2009.

[8] György Szabó, Jeromos Vukov, and Attila Szolnoki. Phase diagrams for an evolutionary prisoner's dilemma game on two-dimensional lattices. *Physical Review E*, 72(4):047107, 2005.

[9] György Szabó and Csaba Tőke. Evolutionary prisoner's dilemma game on a square lattice. *Physical Review E*, 58(1):69, 1998.

[10] Jeromos Vukov, György Szabó, and Attila Szolnoki. Cooperation in the noisy case: Prisoner's dilemma game on two types of regular random graphs. *Physical Review E*, 73(6):067103, 2006.