The Evolution of Cooperation in the Centipede Game with Finite Populations

Rory Smead

January 15, 2007

Abstract

The partial cooperation displayed by subjects in the Centipede game deviates radically from the predictions of traditional game theory. Even standard, infinite population, evolutionary settings have failed to provide an explanation for this behavior. However, recent work in finite population evolutionary models has shown that these settings can produce radically different results from the standard models. This paper examines the evolution of partial cooperation in finite populations. The results reveal a new possible explanation that is not open to the standard models, and thereby gives us reason to be cautious when employing these otherwise helpful idealizations.

1 Introduction

The Centipede game, first introduced by Rosenthal (1981), provides a setting where rational self-interest conflicts with socially optimal cooperative behavior. Two players take turns at an opportunity to acquire the larger share of a sum of money or pass the choice to the other and thereby increase the sum. The game ends when a player takes the money or after a set number of turns, in which case the money is divided equally. Several of the most popular solution concepts in traditional game theory tell us that we should expect to see absolutely no cooperation (passing) in the Centipede game. However, experimental results deviate wildly from these expectations, subjects seem to invariably exhibit partial (though, not total) cooperation.

Rosenthal (1981) initially argued that we should treat Centipede-like situations as pair of single person decision problems. And, from a decision-theoretic standpoint, if each player in the centipede game has certain views about the propensity of the other to deviate from the game-theoretic predictions, then it will be rational to pass in early rounds (96-97). But why should each expect the other to deviate from the game-theoretic solution? We could augment Rosenthal's analysis with the hypothesis that Humans simply have a natural propensity to cooperate, at least partially, in situations like the Centipede game. This leaves us with a puzzle: how can explain this propensity?

For such a question, it is natural to turn to evolutionary game theory. However, the standard evolutionary models do not provide an answer, yielding the same solution as the traditional game-theoretic concepts and so do not aid us in explaining the existence of partial cooperation. However, the standard evolutionary setting assumes an infinite population and recent studies into finite population evolutionary dynamics have revealed some important differences between these finite cases and their infinite counterpart. For instance, Imhof, Fudenberg and Nowak (2005) show that stochastic evolution in finite populations need not select for a strictly dominant strategy. And Taylor, Fudenberg, Sasaki, and Nowak (2004) show that this evolutionary dynamic can even select *against* such strategies.

I will examine the evolution of partial cooperation in finite populations using the

Moran process (Moran 1962) with frequency dependent fitness. Through the use of simulations, I find that, in finite populations, evolutionary paths often "get stuck" in states where the population is exhibiting partial cooperation rather than the unique stable point of no cooperation in the infinite setting. This result illustrates a new potential explanation for the evolution of partial cooperation as well as providing a striking example of how finite population evolutionary models can differ radically from their infinite counterparts. The frequently used and mathematically useful idealizations of the infinite population setting can lead us away from viable evolutionary explanations of partial cooperation. More generally, this suggests that one ought to be cautious when relying on such idealizations in evolutionary models.

I will begin by setting out two games of partial cooperation: the Centipede game and the Quasi-Centipede. In section 3, I will examine these games in the standard evolutionary setting using the traditional replicator dynamic. And, in section 4, I will explore a model of evolution in finite populations and the effect of this finite setting on partial cooperation.

2 The Centipede and Quasi-Centipede Games

As mentioned above, the Centipede Game is a simple case where self-interest can interfere with more efficient, cooperative behavior. Imagine two players sitting at a table with a sum of money α between them, say \$2. They take turns choosing to either pass (p) or take a significant portion leaving the rest to the other player (t). This continues for a fixed, and known, number of rounds or until someone takes the money. If a player chooses to pass, the sum increases by \$2, if a player chooses to take the money they receive $\beta = \frac{1}{2}\alpha + \$1$ and the other player receives the remaining money $(\alpha - \beta)$. If no player chooses to take the money, it is divided as if the player who would have gone next takes the money. The payoffs are such that if the game ends immediately after a player passes, then that player would have been better off taking the money and ending the game.

In the classical example, the game continues for one hundred rounds, hence the name. However, any game with the basic structure and payoff ordering is a Centipede Game. Figure 1 shows a 6-stage version of this game.

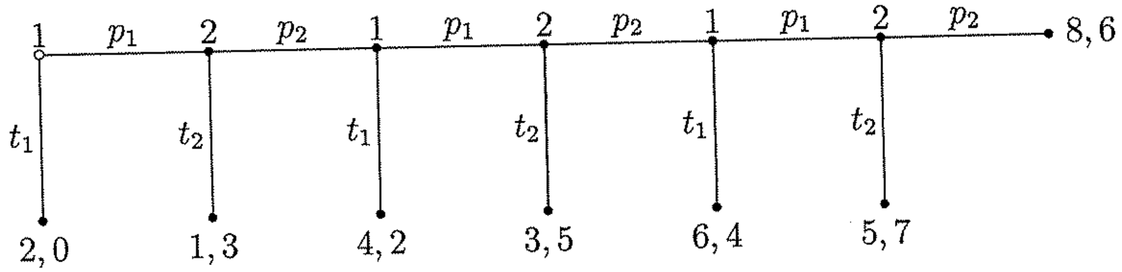

Figure 1: A 6-Stage Centipede Game

This can be easily generalized. Let the strategy sets be $S_1 = \{1, 3, 5, ..., N+1\}$ and $S_2 = \{2, 4, 6, ..., N+1\}$ where N is the length of the Centipede game.[1] And, let $s_i \in S_i$ denote the strategy chosen by player i, this number corresponds to the round on which the player decides to defect (if $s_i = N + 1$, then player i passes at every stage). Suppose $s_i = s_j - 1$ then the utility function for player i is such that $\pi_i(s_i + 2, s_j) < \pi_i(s_i, s_j)$. In other words, if my opponent is going to defect on the

---

[1] This is actually a reduced strategy set relative to the traditional Centipede game. Each strategy should include what the player would choose at every choice point, whether or not that choice point is actually reached. However, for our purposes in this paper this reduced strategy set will suffice.

next round, I am better off defecting now rather than passing. The payoff functions from the above 6-stage Centipede game can be generalized as:

$$\pi_1(s_1, s_2) = \begin{cases} s_1 + 1 & \text{if } s_1 \leq s_2; \\ s_2 - 1 & \text{if } s_1 > s_2. \end{cases}$$

$$\pi_2(s_1, s_2) = \begin{cases} s_1 - 1 & \text{if } s_1 \leq s_2; \\ s_2 + 1 & \text{if } s_1 > s_2. \end{cases}$$

The Centipede Game is often used in game theory courses to introduce the concept of *backward induction* and the *iterated elimination of weakly dominated strategies.* Suppose player two were to find herself at the last stage in the Centipede game above; naturally she would prefer to take the money. But, given that player two would take the money at the last stage, player one should take the money on the second-to-last stage and given player one's preference player two should defect on the third-to-last stage. Iterating this reasoning leads to the conclusion that player one should take the money on the first stage. The unique *subgame-perfect equilibrium* of this game is for each player to opt out at any chance she gets. Every *Nash equilibrium* in this game will share the feature of first-stage defection with the subgame-perfect equilibrium.[2] It is this uniquely "rational" (and unintuitive) solution that traditional game-theoretic wisdom tells us to expect.

---

[2] There are actually a large number of Nash equilibria in the Centipede game. Any pair of (mixed) strategies where player 1 defects immediately and player 2 would choose to defect on her first choice frequently enough to make player 1's immediate defection is a best-response will qualify as a Nash equilibrium. However, the only pure-strategy Nash Equilibria is where $s_1 = 1$ and $s_2 = 2$.

Given that the payoffs for a few rounds of cooperation (passing) are so much better than just taking the money immediately, it seems unlikely that people would actually conform to the "rational" solution. And indeed, in experimental settings, subjects do tend to show some degree of cooperation. McKelvey and Palfrey (1992) studied the behavior of subjects in both 4-stage and 6-stage centipede games at varying stakes. They find that subjects display a substantial amount of cooperation, typically making it more than half-way down the centipede before one takes the money (808).[3] They also found that higher stakes slightly lowered the degree of cooperation, but by a relatively small amount.[4] It is this partial cooperation that stands in need of explanation.

The fact that people regularly deviate from the predictions of traditional game theory has led many scholars to look for alternative explanations. For instance, McKelvey and Palfrey show that if we assume that some individuals are altruists (unconditional passers), then it may be within the scope of "rational" self-interest to play pass. McKelvey and Palfrey also look at models with errors in action or belief and show that partial cooperation could also arise in these settings. In this theme, Binmore (1994) suggests that Centipede game shows us that we need to reassess what constitutes a "rational" player. Nagel and Tang (1998) look to various models of learning with the hope of providing some insight. They find that a model of simple reinforcement learning does fairly well in predicting partial cooperation.

---

[3] One difference between the Centipede game in Figure 1 and the Centipede game used by McKelvey and Palfrey was that the payoffs increased exponentially. The results that will be presented in this paper will also hold for these exponentially increasing Centipedes.

[4] For a survey of various experimental results on dominance-solvable games like the Centipede game see Camerer (2003, 218-236).

These explanations, however, leave open the evolutionary puzzle of why we should see partial cooperation (or altruism) in the first place; as will be seen below, even mindless replicators will converge on the backwards induction solution.

For the sake of illustration the model that will be investigated below will largely focus on a simultaneous move and symmetric game similar to the Centipede game (which I will call a "Quasi-Centipede game").[5] This game begins with a sum of money between the players and at each round they simultaneously declare to take or pass. If both players choose 'take' then they split the money and the game is over. If one chooses 'take' and the other 'pass' the taker receives the majority, leaving the rest to the other and ending the game. If they both choose 'pass' then the sum increases and they play again. This continues until one or both chooses 'take' or for a finite and known number of rounds, at which point the money is split equally. Figure 2 shows a 3-Stage Quasi-Centipede game; the possible strategies are simply the number of rounds k that each player chooses to "pass."[6]

The payoff function in this game can easily be generalized to the N-stage game,

---

[5] This modified game simplifies the model significantly and, as will be shown, the qualitative results will hold for the standard Centipede game as well. Also, experimental results of subjects playing a game that is somewhat similar to the one above closely resemble those of the standard Centipede game. See Van Huyck, Wildenthal, and Battalio (2001) for these results; the term "Quasi-Centipede" is used by Camerer (2003).

[6] To further motivate the connection between the Quasi-Centipede game and the standard Centipede, note that every Quasi-Centipede game, in terms of expected payoffs, may be translated into a scenario where a player is unsure of his role (whether player one or two, each with probability $\frac{1}{2}$) in a standard Centipede game and must simply decide how many rounds to pass before defecting.

		Player 2		
		$s_2 = 0$	$s_2 = 1$	$s_2 = 2$
Player 1	$s_1 = 0$	1, 1	4, 0	4, 0
	$s_1 = 1$	0, 4	3, 3	6, 2
	$s_1 = 2$	0, 4	2, 6	5, 5

Figure 2: A 3-Stage Quasi-Centipede Game

here is one example:

$$\pi_1(s_1, s_2) = \begin{cases} 2(s_1 + 1) + 2 & \text{if } i < j; \\ 2s_1 + 1 & \text{if } i = j; \\ 2(s_2 + 1) - 2 & \text{if } i > j. \end{cases}$$

Where $0 \leq s_1, s_2 < N$ are the player's strategies. The payoff function for player two is similar. This payoff function will be used in the model examined below.

Once again, the unique solution for this game is to take the money immediately, and this solution can be delivered by an iterated elimination of weakly dominated strategies. But, the force of this solution is even stronger than in the Centipede game since there is a unique Nash Equilibrium at $(s_1 = 0, s_2 = 0)$. Thus, in some sense, cooperation in the Quasi-Centipede game may be even more difficult to achieve than in the standard Centipede game.

3 Evolution in the Centipede Game

For the evolutionary setting, we imagine an infinite population of players that are "hard-wired" to cooperate for a certain number of rounds in the Centipede game.

Individuals are paired randomly to play the game and receive payoffs which function as the "fitness" of using a particular strategy k. Let $0 \leq x_s \leq 1$ represent the frequency of strategy s in the population ($\sum_{s \in S} x_s = 1$) and $X = (s_0, s_1 ...)$ represent the state of the population. The fitness of type $s' \in S$ (the set of individuals who use strategy s') is calculated by:

$$f_{s'}(X) = \sum_{s \in S} \pi(s', s) x_s$$

Let θ_X be the average fitness of the population in state X; the distribution of strategies evolves according to the standard replicator dynamic.[7] The differential equation governing the change in frequency of type s is:

$$\dot{x_s} = x_s(f_s(X) - \theta_X)$$

In words, types that have higher fitness increase in proportion at the expense of those with lower fitness.

The replicator dynamic does not generally support the iterated elimination of weakly dominated strategies, a fact that can lead to interesting results. Thus, there may be some hope for promising results in the Centipede game using the replicator dynamic. However, in both the Centipede game and Quasi-Centipede game the replicator dynamic gives us a single result: a population of first-round defectors.

In the case of the 3-stage Quasi-Centipede game above, it is easy to see that there is a unique evolutionary outcome: everyone plays 0. If we suppose the population

---

[7] For details on this dynamic, see Hofbauer and Sigmund (1998).

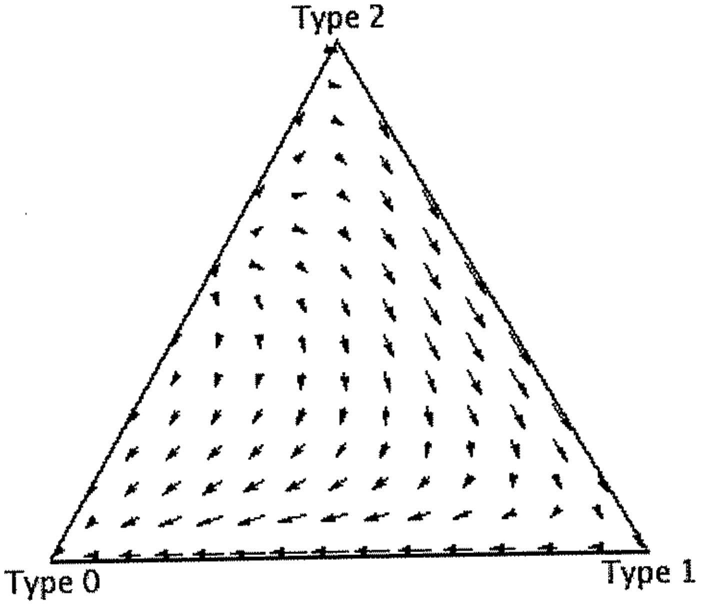

Figure 3: Simplex of a 3-Stage Quasi-Centipede Game

consists almost entirely of type 2, it will be invaded by individuals of type 1 and if the population consists of almost entirely type 1 then it will be invaded by individuals of type 0.[8] The same result holds for arbitrarily large Quasi-Centipede games. Figure 3 shows the global dynamics for a 3-stage Quasi-Centipede game. It is interesting to note that the dynamics do not necessarily carry the population directly to the unique equilibrium. Instead, the dynamic sometimes seems carry the population through the stages of an elimination of dominated strategies, first to a predominance of type 2, then to type 1, and finally to the evolutionary equilibrium at all type 0.

As seen in figure 3, some paths of evolution in this 3-stage game come close to a non-equilibrium edge of the simplex. The longer the N-stage Quasi-Centipede game,

---

[8] This is under the assumption that every type is represented in the initial population. The replicator dynamics will then only carry some types to extinction in the limit, a point that will be important when we turn to finite populations.

the more likely it is to have an evolutionary path that gets close to a non-equilibrium edge of the simplex. The reason for this is that there is no direct route from type $s_i = k$ to type $s_i = k - 2$ (for any $N \geq k \geq 2$) and so, to evolve the unique solution as mentioned above, the population must evolve through the stages of backwards induction.

The evolutionary results for the standard Centipede game are similar to the Quasi-Centipede game.[9] In a sufficiently mixed population, types that show a high-degree of cooperation initially outperform those that do not and hence, increase in frequency. Then, those types that exhibit just slightly less cooperation grow at the expense of the more cooperative type; and this process repeats until the population consists only of first-round defectors. In general, stable states of the replicator dynamic must also be Nash equilibria and all Nash equilibria in the Centipede game have the property of first-round defection.[10] Giovanni Ponti (2000) shows that this fact holds for a large class of dynamics that includes the replicator dynamic; proving that these dynamics will eventually carry the population to a state that is "outcome equivalent" to the sub-game perfect equilibrium (total first-round defection). Ponti also notes that populations often exhibit "unlearning" of the solution (increasing in cooperation) before converging on first-round defection and he shows that as the length of the Centipede increases, the more likely populations are to exhibit

---

[9] The possible types have to be adjusted for the Centipede game. Specifically, an individuals strategy must include what the individual should do if in the role of player 1 and what to do if in the role of player 2.

[10] Some weakly dominated strategies will survive in small amounts. In a population of first-round defectors, there is no difference in payoff between strategies with regard to the action when in the role of player 2. Thus, some strategies which are cooperative when in the role of player 2 may survive in stable states, but there is always unanimous first-round defection when in the role of player 1.

"unlearning."

At this point we can conclude that the standard evolutionary models (the replicator dynamic) cannot alone account for the evolution of partial cooperation. Even mindless replicators can evolve the backwards-induction solution to the Centipede game. However, if we include some perturbation in the dynamics (which could be interpreted as either mutation or experimentation), we no longer see the first-round defection result. In addition to his results mentioned above, Ponti also shows that with enough perturbation in the dynamic a population can exhibit cycles between higher and lower degrees of cooperation.

The cyclic dynamical picture given by mutation in the Centipede game provides one potential explanation for the partial cooperation seen in experiments. In the following section I will examine a model that will provide another potential evolutionary explanation that does not rely on mutation.

4 Partial Cooperation in Finite Populations

In recent years, Taylor, Fudenberg, Sasaki, and Nowak (2004) have developed a stochastic model of evolution in finite populations by incorporating frequency dependent fitness into the Moran process (Moran 1962).[11] Taylor et. al. have shown that this process can deliver radically different results than the standard replicator dynamic. In this section, I will use the Moran process to model the evolution of partial cooperation in the Centipede and Quasi-Centipede games. I find that without

---

[11] This process has also been studied in a series of papers by Fudenberg, Imhof, Nowak and Taylor (2004, 2006).

any mutation (or a very small amount), resulting populations will regularly exhibit partial cooperation.

On this model, the population consists of M individuals. Let $x_s^t \leq M$ denote the number of individuals of type s at time t, then $\sum_{s \in S} x_s^t = M$. And let $X_t = (x_{s_1}^t, x_{s_2}^t, \ldots)$ represent the state of the population at time t. The fitness to a given type s' at time t in the population is:

$$f_{s'}^t(X_t) = \pi(s', s')(x_{s'}^t - 1) + \sum_{s \neq s'} \pi(s', s) x_s^t$$

The evolution is stochastic. At each step, one individual is chosen to reproduce with a probability proportional to her fitness and produces one identical offspring. Then, another individual is selected randomly to die. The probability for reproduction of type s' at time t is:

$$R_{s'}^t(X_t) = \frac{x_{s'}^t f_{s'}^t}{\sum_s x_s^t f_s^t}$$

And the probability for a death of a type s is simply $\frac{x_s^t}{M}$. We can also talk of selection favoring certain strategies at certain states of the population, we can say that selection favors x_s in state X_t if $R_s^t(X_t) > R_{s'}^t(X_t)$ for all $s' \neq s$.

As the population size becomes very large ($M \to \infty$) and time is appropriately adjusted, this process approximates the standard replicator dynamic (Fudenberg et. al. 2004, 2006). But for finite populations (without mutation),[12] a population governed by the Moran process is effectively a large finite-state Markov chain with as many absorbing states as there are types in the population. Once $x_s^t = M$ for

---

[12] The effects of mutation in these models will be considered later in this section.

some s, that type has completely overtaken the population and becomes the only type that can reproduce ($R_s^t = 1$), hence the population stays in this state. Also, if the frequency x_s^t of some type s drops to $x_s^t = 0$, then it cannot increase since $R_s^t = 0$ and so extinction of a type is possible on this dynamic. But what happens to partial cooperation in this evolutionary setting?

4.1 The Quasi-Centipede in Finite Populations

To answer this question, I will first focus on the Quasi-Centipede game and then turn to the Centipede game. I will assume that the Quasi-Centipede game being played is the same as the generalized one mentioned above and will employ the use of simulations. Just as in the traditional evolutionary setting, we will not necessarily get a direct trajectory to the evolutionary equilibrium, instead an appropriately mixed population will take a round-about path similar to some shown in Figure 3. However, under the Moran process we see several differences. First, the dynamic is discrete-time rather than continuous. Second, the step-direction is stochastic, only approximating the force of selection on average. Third, as a result of these differences, it is almost always possible for a type to go extinct in this setting.

The fact that extinction is possible and that the dynamic is stochastic can now cause "breaks" in the evolutionary chain that follows the iterated elimination of dominated strategies. This can result in the population being absorbed into a state where every member is playing a partially cooperative strategy. Recall Figure 3, in this setting we can imagine an evolutionary path close to the Type 2 - Type 1 edge of the simplex where Type 0 happens to go extinct.[13] In this case, the population will most likely be absorbed into a state where every individual is playing Type 1 (since Type 1 out performs Type 2 in this population). If extinction of strategies in the Quasi-Centipede game is possible, then a population can "get stuck" along its evolutionary path and end up playing a partially cooperative strategy. But, how likely is such an outcome?

The answer to this question depends on many factors: the size of the population, the length of the Quasi-Centipede, and the initial state of the population. I will begin with simulations in an example case. Suppose $M = 1000$ and suppose that each individual is given a random strategy, with equal probability for each strategy (equal initial weights), for a 10-Stage Centipede game, simulations show that an overwhelming likelihood of resulting populations ending in an absorbing state with a substantial degree of cooperation. In 1000 trials, there was not a single case of the unique equilibrium seen in the traditional setting and most ended in partial cooperation over the half-way point in the Centipede in an average of just over 150,000 generations.

One may object that this method of assigning starting strategies is overly biased toward the center of the population's state-space. This can be remedied by selecting a random point on the simplex of strategies and assigning the weights accordingly (random initial weights).[14] This modification causes more variance in terms of which absorbing state is reached, but we still observe partial cooperation with very high

---

[13] For instance, there could be only one Type 0 in a population of Type 1's and 2's, Type 1 could be chosen to reproduce and the lone Type 0 could be chosen to die leaving $x_0 = 0$.

[14] This is done by assigning each strategy s a random weight on the interval [0,1] w_s and assigning an individual strategy s with probability $\frac{-ln(w_s)}{\sum_{s' \in S} -ln(w_{s'})}$.

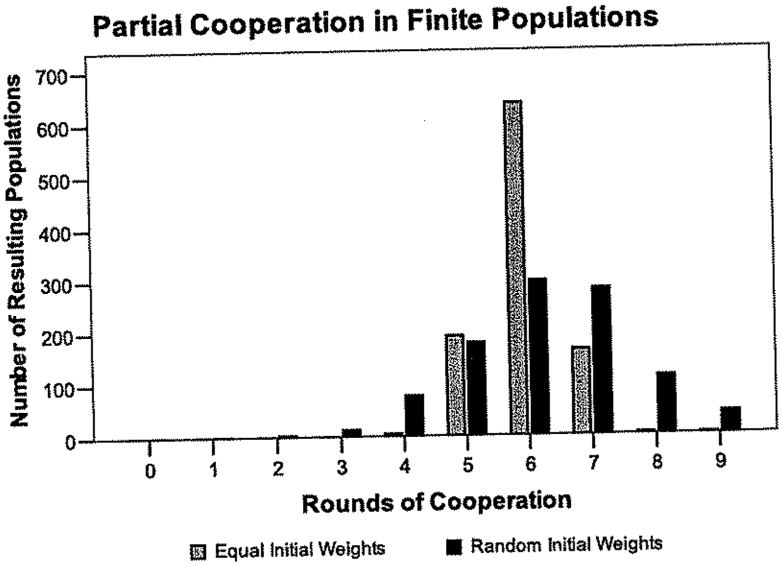

Figure 4: The results of 1000 simulations on a 10-stage Quasi-Centipede game with a population of 1000 evolving by the frequency dependent Moran process.

frequency. Figure 4 gives a summary of the results for 1000 trials under both styles of random initial conditions. In each case, we can see that it is very likely (probability greater than 0.99) that an evolving population will become absorbed into a state where every member exhibits partial cooperation.

As mentioned above there are other key variables that affect the resulting degree of cooperation in a population: the length of the Centipede and the size of the population.[15] If we increase the population size M then the average level of cooperation seen in the resulting population decreases. The reason is that the step-size $(\frac{1}{M})$ of the evolution becomes proportionally smaller as M increases, and so the population

---

[15] The specific values of the payoffs also make a difference in the outcome. We can, while preserving the ordinal raking of payoffs, make things harder or easier for cooperation to result by making the payoff for defecting larger or smaller. However, the qualitative results are robust across such changes.

is less likely to get extremely close to the edge of the simplex.

If, on the other hand, the size N of the Quasi-Centipede is increased, then we see the average level of cooperation in the resulting population increase. The reason is that, as N increases the number of places extinction of strategies could occur increases. Since the evolution roughly follows the stages of iterated elimination of weakly dominated strategies there will be $N-1$ edges of the simplex that correspond to steps in the elimination of dominated strategies. And, each time the evolution comes near an edge of the simplex, there is a chance of strategy extinction, and hence, a chance of becoming absorbed into a partially cooperative state.

Figure 5 shows the average level of cooperation (normalized on $[0,1]$ where 0 corresponds to no cooperation and 1 to complete cooperation) as a function of the population on a logarithmic scale for 5, 7, 10, and 14 stage Quasi-Centipede games. These simulations show that the larger the population, the more the dynamics resembles the infinite case, as expected. And, the longer the Centipede the higher the expected degree of cooperation in the resulting population.[16] Given these results, it is reasonable to conjecture that for any finite population size M there will be some N-stage Centipede game (with payoffs as specified above) that will result in partial cooperation with a probability arbitrarily close to 1.

---

[16] With very short Quasi-Centipede games (such as $N = 3$) and the payoffs as in Figure 2, populations resulting in partial cooperation are rarely observed. However, increasing the potential benefit for cooperation can increase the frequency of partial cooperation.

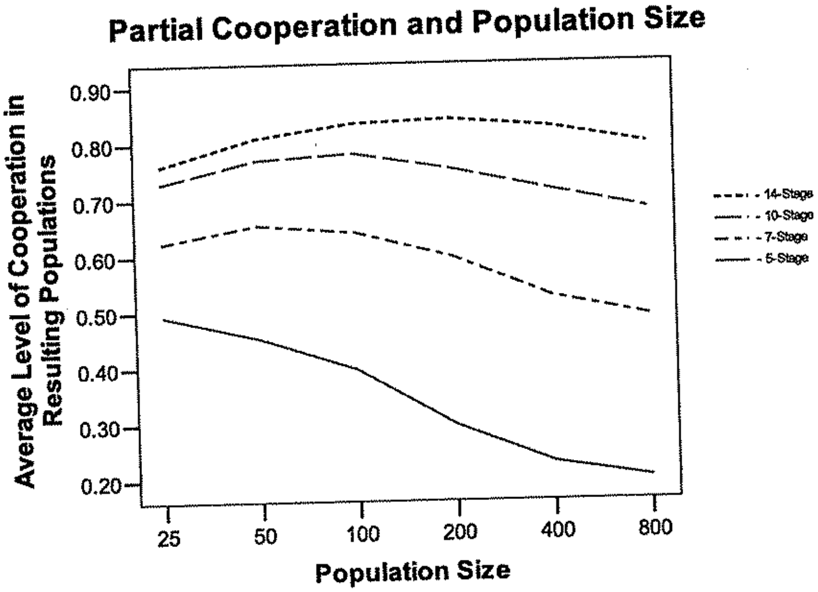

Figure 5: The average resulting levels of cooperation in 5, 7, 10, and 14-stage Quasi-Centipede games as a function of population size.

4.2 The Centipede game in Finite Populations

Thus far, I have focused on the Quasi-Centipede game. Simulations on the Centipede game using the Moran process show that the same qualitative results hold for this game. For these results, we imagine two populations playing the Centipede game specified in section 2, where members from one always take the role of player 1 and members from the other always take the role of player 2. The fitness for a type in a population is calculated by how well they do against the other population and the evolution within a population is governed by the Moran process.[17]

---

[17] The single population setting simply has more types, where an individual's type must specify what to do in the role of player 1 and the role of player 2. The results for the single population are similar to the two-population case. However, single populations often persist in awkward polymorphic mixes for prolonged periods of time before reaching an absorbing state. In these mixes, every member defects on some round n when in the role of player 1 but there is no consensus on

In an 8-stage Centipede game being played between 2 populations of $M = 1000$, as in the case of the Quasi-Centipede, there was no resulting populations of first-round defectors in 1000 trials.[18] And, the average level of cooperation was over the half-way point of the Centipede (the average node that resulting populations played "take" was 6.46) with the average time to an absorbing state was just under 200,000 generations. We also find that the same general qualitative results hold as in the Quasi-Centipede game, the average level of cooperation in resulting populations increases with the length of the Centipede and decreases as population size M increases. And, as before, these changes are due to the changes in likelihood of strategies going extinct.

Thus, in both the Centipede and Quasi-Centipede games, the Moran process illuminates a new possibility for the evolution of partial cooperation. Since the evolution tends to follow the stages of elimination of dominated strategies, it is easy for certain strategies to become extinct. This can cause the population to become absorbed into a partially cooperative state, and these simulations have shown it is very likely to do so.

4.3 Mutation

Although the Moran process is stochastic, the models examined so far have not included any mutation. Indeed, the lack of mutation plays a key role in finite popula-

---

which round $m > n$ to defect on when in the role of player 2. For this reason, only results of the two-population case are presented here.

[18] The initial conditions for these simulations were determined by assigning each individual a random strategy with equal probability across strategies. If strategies were assigned with random weights, as above, the same qualitative results seen in the Quasi-Centipede game hold: more variety in resulting populations is observed, but an overwhelming majority show high degrees partial cooperation.

tions resulting in partially cooperative states since this result relies on the extinction of certain strategies. We can introduce a mutation probability μ into the reproduction process, such that, with probability μ a new individual picks a random strategy (with equal weights to each strategy) instead of keeping the same strategy as the reproducing type.

By introducing a very small mutation rate, the results presented above break down in the limit (as $t \to \infty$). In this setting, the population settles into a normal partially-cooperative state but eventually just the right mutation(s) will occur to cause a slightly less cooperative type to take over. However, for these very small mutation rates, it may take a very long time to reach the immediate-defection state. For instance, simulations in a 10-stage Quasi-Centipede game with a population of $M = 1000$ and a mutation chance of 1-in-10,000, can take several million generations before reaching the immediate-defection state. Furthermore, if the mutation rate is high enough (say, 0.02 in the 10-stage, $M = 1000$ case), then we see a finite-population analogue to Ponti's (2000) cycles of evolution seen in infinite populations with perturbation. Figure 6 shows the average level of cooperation over time in such a cycling population through the first 1-million generations.

Similar evolutionary cycles have been seen in other games studied with the Moran process in finite populations. Imhof, Fudenberg, and Nowak (2005) have examined the repeated prisoner's dilemma in this setting and have shown that mutation can cause cycles between cooperative, uncooperative, and discriminatory strategies. In the Centipede and Quasi-Centipede games, there is a window of mutation that will cause partial cooperation to break down and result in populations arriving at the

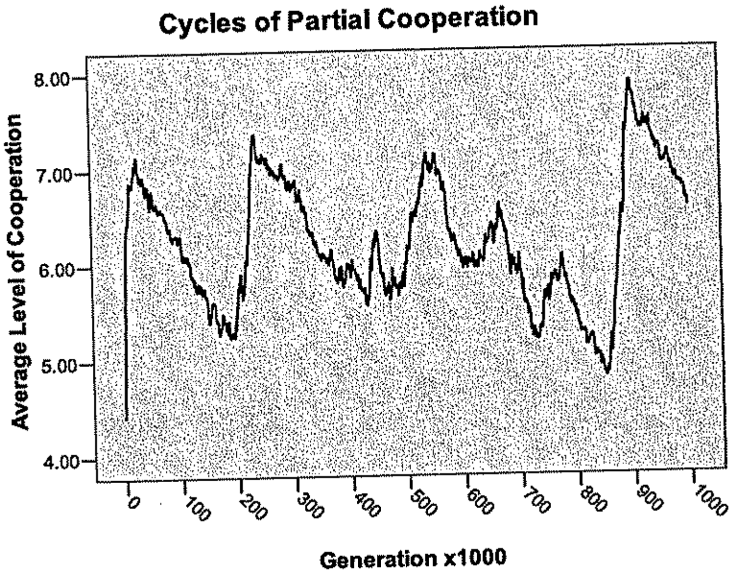

Figure 6: Cycles of partial cooperation in an evolving population of 1000 individuals playing a 10-stage Quasi-Centipede game with a mutation rate of $\mu = 0.02$.

immediately-defect solution (mimicking the infinite population setting without mutation). In this window, cycles are possible, but rarely observed in simulations. If $\mu = 0$ or is very small, then the population will get stuck in states of partial cooperation indefinitely or for very long periods of time. If μ is large, then we see frequent and indefinite cycles of partial cooperation in the population.[19]

5 Conclusion

The move to finite populations with the possibility of extinction can have a radical effect on the evolution of behaviors in certain games. We have seen that this is

---

[19] The exact place and size of this window will vary with population size and the length of the Centipede.

certainly the case with the Centipede and Quasi-Centipede games. In these games, the evolution roughly follows the stages of iterated elimination of weakly dominated strategies and in the standard infinite population models results in no cooperation. However, because of the nature of finite population evolutionary dynamics, certain strategies can go extinct causing the population to "get stuck" in partially cooperative states. The likelihood of establishing partial cooperation increases with the size of the game and decreases with the size of the population.

This result illustrates a new potential explanation for the evolution of partial cooperation: it may be that evolutionary paths, although leading in the direction non-cooperative states, become absorbed into partially cooperative states due to the finite and stochastic nature of the evolutionary process. Additionally, this result provides a striking example of how evolutionary processes in finite settings can deviate significantly from those in the traditional infinite setting. The idealizations of the latter render such evolutionary outcomes as seen in finite cases impossible, and consequently, blind us to the new potential explanations for partial cooperation.

References

[1] Binmore, K. (1994) "Rationality in the Centipede" *Theoretical aspects of knowledge: Proceedings of the fifth conference* 150-159.

[2] Camerer, C. (2003) *Behavioral Game Theory: Experiments in Strategic Interaction.* New Jersey: Princeton University Press.

[3] Fudenberg, D., Imhof, L., Nowak, M., and Taylor, C. (2004) "Stochastic Evolution as a Generalized Moran Process." working paper.

[4] Fudenberg, D., Nowak, M., Taylor, C., and Imhof, L. (2006) "Evolutionary Game Dynamics in Finite Populations with Strong Selection and Weak Mutation." *Theoretical Population Biology* 70: 352-363.

[5] Hofbauer, J. and Sigmund, K. (1998) *Evolutionary Games and Population Dynamics.* Cambridge University Press.

[6] Imhof, L., Fudenberg, D., and Nowak, M. (2005) "Evolutionary cycles of cooperation and defection." *PNAS* 102: 10797-10800.

[7] Moran, P. A. P. (1962) *The Statistical Processes of Evolutionary Theory,* Oxford: Clarendon Press.

[8] McKelvey, R. and Palfrey, T. (1992) "Experimental Results on the Centipede Game in Normal Form: An Investigation on Learning." *Econometrica* 60: 803-836.

[9] Nagel, R. and Tang, F. (1998) "Learning, Local Interaction, and Coordination." *Journal of Mathematical Psychology* 42: 356-384.

[10] Nowak, M., Sasaki, A., Taylor, C., and Fudenberg, D. (2005) "Emergence of Cooperation and Evolutionary Stability in Finite Populations." *Nature* 428: 646-650.

[11] Ponti, G. (2000) "Cycles of Learning in the Centipede Game." *Games and Economic Behavior* 30: 115-141.

[12] Rosenthal, R. (1981) "Games of Perfect Information, Predatory Pricing, and the Chain Store." *Journal of Economic Theory* 25: 92-100.

[13] Taylor, C., Fudenberg, D., Sasaki, A., and Nowak, M. (2004) "Evolutionary Game Dynamics in Finite Populations." *Bulletin of Mathematical Biology* 66: 1621-1644.

[14] Van Huyck, Wildenthal and Battalio. (2002) "Tacit Cooperation, Strategic Uncertainty, and Coordination Failure: Evidence from Repeated Dominance Solvable Games." *Games and Economic Behavior* 38: 156-175.