

# Tag Mediated Cooperation with Multi-Level Selection

Terence Soule \*  
The University of Idaho

October 16, 2007

## Abstract

The conditions and mechanisms leading to cooperation in evolving systems remains both a fundamental and complex question. Although significant progress has been made in mapping the conditions that support cooperation in simple models, the evolutionary trajectories that lead from non-cooperative conditions to cooperation, particularly in more complex systems, are still poorly understood. In this paper we study how cooperation levels vary in both a simple tag-mediated cooperation model and a more complex multi-level version of the same system in which there is competition between both individuals and populations of individuals. The results confirm that including inter-population competition increases the average level of cooperation within individual populations. However, the mechanisms leading to the increased cooperation are unexpected. First, with multi-level competition, the periods of cooperation within a population become more frequent, but generally remain unstable, and periods of almost pure defection are still common. Thus, the average amount of cooperation is increased, but no individual population maintains a stable level of cooperation. Second, with inter-population competition, when a threshold level of cooperation is reached the entire system is rapidly driven to a state of almost pure cooperation that is stable. This state does not occur when competition is only within populations.

## 1 Introduction

In 2005 Science magazine included the question: ‘how did cooperative behavior evolve?’ as one of the 25 most fundamental questions facing the scientific community [1]. Although cooperative behavior is clearly widespread in both nature and human interactions, altruistic cooperation, in which individuals sacrifice their own fitness to improve some one else’s fitness, runs counter to simple interpretations of evolutionary principles.

---

\*This research supported by NSF Grant #0535130.

Typically the evolution of cooperation depends on specific features that lead to specific mechanisms fostering cooperation. Examples of such mechanisms include kin selection, reciprocity (direct, indirect, and network), and group selection. It is worth noting that these mechanisms are not mutually exclusive. Nowak has outlined the conditions under which each of these mechanisms should lead to cooperation, at least for very simple models [2]. Understanding how cooperation evolves when multiple mechanisms are present remains a difficult problem.

A common case where multiple mechanisms support cooperation is multi-level competition, in which there is competition at both the individual and the population level. In such a model there may be evolutionary pressure in favor of, or opposed to, cooperation at either or both levels. For example, within populations cooperation may benefit individuals through kinship, whereas competition between populations of individuals may favor cooperation within the population independent of kinship.

Riolo et al. [3] and Spector and Klien [4] have shown cooperation can evolve in a simple evolutionary model with tag-mediated cooperation. In this paper we use their tag-mediated system to show that the addition of multi-level competition, competition between individuals and simultaneously competition between populations of those individuals, has several significant effects on the evolution of cooperation. First, in general we find that cooperation levels are unstable, with cooperation rates undergoing rapid and extreme changes. Second, our results confirm that multi-level competition can increase the average cooperation rate by changing the frequency and duration of the periods of cooperation. Finally, and perhaps most significantly, multi-level competition introduces a new stable state in which all of the competing populations have near 100% cooperation. Once this state is reached defection becomes an unsuccessful strategy and the state remains stable. Thus, multi-level competition has both quantitative effects (increasing the frequency and duration of the periods of cooperation) and qualitative effects (introducing a new stable state) on the evolutionary dynamics of cooperation.

## 2 Background

In 2001 Riolo et al. showed that cooperation mediated by tags could arise without reciprocity [3]. Individuals *donate* to other individuals whose tag is not too dissimilar to the donor's tag. Donation involves losing resources and a corresponding decrease in the chance of survival, while the receiver's resources, and chance of survival, increase. In this model (based on a similar model by Holland [5]) there is no reciprocity - the donor gains no direct advantage and the receiver and donor do not automatically exchange roles.

Donations in this model are a form of strong altruism, as donors receive an absolute penalty [6, 7]. In contrast, in *weak* altruism donors receive a benefit, but one that is smaller than received by the other members of the group [6]. Similarly, according to definitions this is a model of an "other-only" trait, as

opposed to a “whole-group” trait in which all members of the group, including the donor, benefit [8].

Initially there was some debate over the interpretation of Riolo et al.’s model because individuals were required to share with other individuals with identical tags [9, 10]. Thus, individuals that were purely selfish could not evolve. In 2006 Spector and Klein showed that even if the restriction that individuals must cooperate with other individuals with identical tags is removed, significant levels of cooperation evolve [4, 11]. They found that the rate of cooperation depends on both the rate of mutation and the size of the ‘neighborhood’ in which individuals both share resources and compete. In particular low mutation rates and small neighborhoods of interaction favor higher average levels of cooperation. Thus, the model allows for adjustable rates of evolved cooperation, making it a very useful empirical model. In these experiments we use a mutation rate of 0.01 and two neighborhood sizes: 5 and 10. From Spector and Klein’s results these parameters are known to lead to significantly different levels of cooperation making them useful for comparing the effects of multi-level cooperation.

Spector and Klein also proposed a likely evolutionary mechanism by which cooperation would appear in a non-cooperative group in this model. However, they did not present the long term evolutionary dynamics of individuals runs. In particular, the higher average levels of cooperation observed with low mutation rates and small neighborhoods could arise from several factors: higher levels of cooperation in all members, very high levels of cooperation by some members and lower levels by other members, or fluctuating levels of cooperation that have higher averages over time.

In addition to questions regarding the dynamics of evolved cooperation there is considerable interest in the role of multi-level selection - selection between individuals within groups and between groups - on the evolution of cooperation. Recently Killingback et al. showed that group competition can increase donation rates in a tragedy of the commons model in which individuals donate to the communal ‘pot’, [12]. This is a model of weak, whole-group altruism - individuals are benefited by their donations, although defectors enjoy a larger net benefit.

Bowles et al. have shown that group competition can increase the frequency of strongly altruistic donations [13]. In their model individuals are either pure altruists (always donate at a cost to themselves) or pure defectors (never donate). They also found that within group structures that increase the probability of altruists interacting with other altruists and defectors interacting with other defectors (what they term segmenting) improves cooperation rates. This is similar to Spector and Klein’s results that show that smaller neighborhoods within groups increases cooperation rates [4].

However, unlike tag mediated models, in both Killingback et al. and Bowles et al.’s models individuals are always either altruists or defectors (in Killingback et al.’s model individuals can vary how much they donate), that is, individuals can not choose whether or not to donate. Recently, it has been shown that the ability to recognize other potential altruists and thus to preferentially donate to them (i.e. the greenbeard effect) can significantly increase altruistic behaviors,

if the genes for altruism and ‘beard color’ are loosely coupled [14]. Thus, the use of tags represents a distinct mechanism for evolving cooperation that may, or may not, be improved through group level competition.

In this paper we more closely examine the evolutionary dynamics of tag mediated cooperation to determine the fundamental cause of the increased levels of cooperation observed by Spector and Klein. Then we examine how these dynamics change when multi-level competition is included. By using a relatively simple model, its possible to understand the specific mechanisms by which inter-group selection affects intra-group evolutionary dynamics.

### 3 Methods

The multi-level evolutionary model consists of a set of separate populations. Competition takes place between the individuals in the populations and the populations compete with each other creating two levels of competition. The basic algorithm is outlined in Figure 3 and the details of the model are as follows.

The evolutionary model uses a set of 50 populations. Each population contains 100 individuals arranged in a ‘ring’. In each iteration there is a *sharing* phase and a *selection* phase based on Riolo et al.’s [3] model (with Spector and Klein’s [4] modification) of tag mediated cooperation. Individuals consist of two real values, a tag and a ‘tolerance’. Initially these are chosen randomly from the range (0,1.0). During the sharing phase each individual has  $D$  opportunities ( $D = 3$  in these experiments) to donate to randomly selected individuals with a neighborhood of  $N$  of the donor. An individual choses to donate if the difference between its tag and the receiver’s tag is less than the donor’s tolerance. If an individual donates then the donor’s resources are decreased by *cost* and the receiver’s resources are increased by *benefit*. In these experiments  $cost = 0.1$  and  $benefit = 1.0$ . In each iteration the individuals start with 0 resources; resources may go negative.

After the sharing phase each population enters the selection phase. Each individual is compared to a randomly selected individual (a competitor) from within its neighborhood. The individual with more resources is mutated and then placed in a new population (in the case of a tie the original individual is kept). Thus each individual has at least one opportunity to reproduce if it has equal to or more resources than its randomly selected competitor. In addition, an individual has the opportunity to reproduce if it is selected as a competitor. Once the new population is filled it becomes the ‘current’ population and the old population is discarded. The resources of each individual are reset to zero at the beginning of each iteration and are not maintained from iteration to iteration.

During mutation there is a 0.1 chance that a tag will be mutated and a separate 0.1 chance that the tolerance will be mutated. A tag is mutated by randomly choosing a new tag value from the range (0,1.0). A tolerance value is mutated by adding Gaussian noise with mean 0 and standard deviation 0.01. Thus, tolerances can (slowly) increase indefinitely, although the maximum possible tag difference is 1.0 so any tolerance of 1.0 or greater represents an individual

```

repeat 50 times
  generate a random population of 100 individuals
repeat 440,000 times
  for each population
    for each individual within the population
      set the individual's resources to zero
    for each individual within the population
      randomly pick 3 neighbors to (potentially) share with
    for each individual within the population
      randomly select one neighbor to compete with
      copy the winner into a new (next generation) population
    delete the old population
    make the new populations the current population
  repeat C times
    randomly select 2 populations to compete with each other
    replace the loser with a copy of the winner

```

Figure 1: The basic multi-level algorithm. The numbers correspond to the parameters used in this paper, other parameters may be chosen.

who will donate to anyone. As in Spector and Klein's model a minimum tolerance of  $-0.000001$  is imposed, if mutation generates a tolerance less than this the tolerance is reset to  $-0.000001$ . Note that an individual with a negative tolerance will never donate to any other individual, including one with an identical tag.

In addition to competition between individuals within each population there is competition between populations. After every population has undergone one iteration of sharing and selection selection is applied between populations. Two populations (out of the 50) are selected randomly. If the sum of all resources of all individuals in the second population exceeds the sum of all resources of all individuals in the first population, then the first population is replaced with a duplicate of the second population. Note that this is a one way competition; only the population selected second can replace the first population. Although later the populations may be selected in the reverse order, thereby switching their roles in selection.

In general, a population in which the individuals often donate to each other will have more total resources and will replace a population in which individuals rarely donate. The selection process is repeated  $C$  times. Experiments are run with values of  $C$  between zero (no population level selection) and five (five pairs of populations are selected to compete in each iteration). As  $C$  increases there is more pressure for populations to cooperate - populations with low levels of cooperation are more likely to be involved in a competition and replaced by a population with a higher average rate of cooperation.

Note that this is a fairly extreme form of competition. The losing population

Population size	100
Number of populations	50
Donation opportunities per iteration	3
Iterations	440,000
Cost	0.1
Benefit	1.0
Mutation rate	0.1
Neighborhood	5 and 10
Number of second level competitions	0, 1, 2, 3, 4, 5
Number of trials	50

Table 1: Summary of the model parameters

(consisting of 100 individuals) is completely replaced by a copy of the winning population. This model most closely mimics situations in which the individuals can learn, i.e. where a losing population learns and hence adopts the winner’s cooperative strategy and is similar to that used by Bowles et al. [13].

The program is run for 440,000 iterations, where each iteration consists of a sharing and selection phase for each population and  $C$  competitions between populations (as described above). The first 40,000 iterations are discarded to avoid including any transient behavior in the results. Note that iterations are numbered from the first recorded iteration. Fifty independent trials are run for each of the test conditions.

## 4 Results

In the experiments we examine three progressively finer views of the evolutionary dynamics. First, we examine the general behavior averaging across all of the competing populations. This provides the average cooperation rates as a function of neighborhood size and rates on inter-population competition. Next we examine the dynamics within the evolving populations to determine how inter-population competitions affect intra-population dynamics. Finally, as a control, we remove the tags to determine whether the tags actually play a fundamental role in the evolution of cooperation in this model.

### 4.1 Inter-population Dynamics

Figure 2 shows the average cooperation rates of the populations for neighborhoods of size 5 and 10 with varying amounts of population level selection (error bars are one standard deviation). The average rate of cooperation is determined by dividing the number of actual cases in which individuals donate by the number of opportunities individuals have to donate. In these experiments the population size is 100 and in each sharing phase each individual has 3 opportunities to donate. Thus, the average cooperation rate is the number of times

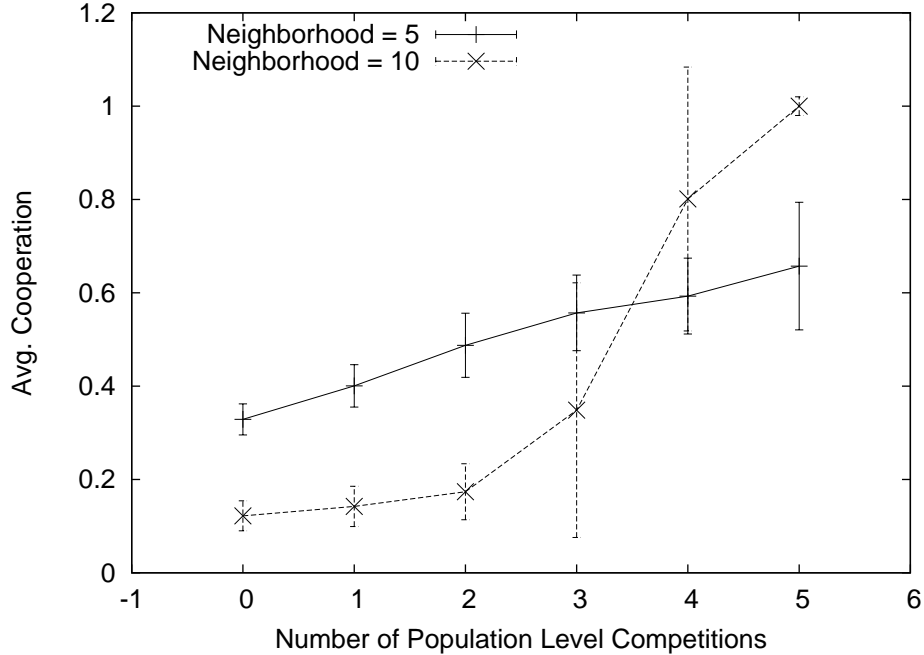


Figure 2: Average cooperation rates for two neighborhood sizes (5 and 10) and several rates of population level competition (0, 1, 2, 3, 4, and 5 population level competitions per iteration). Competition between populations clearly encourages cooperation within populations. Interestingly, this effect is more pronounced as the neighborhood increases.

individuals do share during the sharing phase divided by 300 (100 individuals x 3 opportunities).

When the number of competitions is zero ( $C = 0$ ) there is no selective pressure on the populations. In this case the results agree with those of Spector and Klein [4, 11]. Specifically, the smaller neighborhood results in higher levels of cooperation. As  $C$  increases the average level of cooperation clearly increases demonstrating that the population level competition does increase cooperation within the populations. Interestingly, as the amount of competition between populations increases the neighborhood trend reverses and the larger neighborhood produces more cooperation.

To further illustrate the dynamics of the system, Figure 3 shows the average cooperation rate for 5 sample trials (out of the full 50 trials) with a neighborhood of 10 and 5 population competitions per iteration. Each curve shows the cooperation rate averaged across the 50 populations within the trial. Note that the data does not show anything resembling a constant cooperation rate, instead the average cooperation rate observed in Figure 2 is the result of extreme fluctuations in cooperation rates. These fluctuations continue until a critical




Figure 3: Average cooperation across all 100 populations in 5 sample trials. Neighborhood size is 10 and there are 5 population level competitions per iteration ( $C = 5$ ). In each trial average cooperation across all populations fluctuates until a critical value is reached. Once the critical value is reached the populations quickly move to a stable state with nearly 100% cooperation.

value of roughly 80% cooperation is reached. Once this critical value is reached average cooperation quickly reaches near 100% cooperation. This is a steady state. Once reached significant deviations from 100% cooperation were never observed. This behavior is observed in all five of these sample trials and for these experimental conditions occurred to all 50 trials within the 440,000 iterations of the experiment.

## 4.2 Intra-Population Dynamics

To understand the dynamics of a single population, Figures 4, ??, and ?? show the cooperation rate within a single population compared to the average cooperation rate across all 50 populations in a single, sample trial. (Recall that each of the trials examines 50 competing populations and each population consists of 100 competing individuals.) The parameters used to generate the data in these figures are a neighborhood of 10 with 0, 3, or 5 population level competitions per iteration.

In all three cases, the cooperation rate within a population fluctuates widely






Figure 4: Average cooperation rate for a single population averaged across all 50 competing populations for neighborhood of 10 and  $C = 0$  (no second level competition). The population shows rapidly fluctuating cooperation levels, but the average cooperation level remains low showing that at any given moment only a few (at most) of the populations have a high cooperation rate.




Figure 5: Average cooperation rate for a single population averaged across all 50 competing populations for a neighborhood of 10 and  $C = 3$ . The population shows rapidly fluctuating cooperation levels. The average cooperation level generally remains low showing that generally only a few of the populations have a high cooperation rate. Occasional periods of high average cooperation show that many of the populations have high cooperation rates simultaneously.




Figure 6: Average cooperation rate for a single population averaged across all 50 competing populations for a neighborhood of 10 and  $C = 5$ . Again the single population shows rapidly fluctuating cooperation levels, but the typical level is relatively high. Eventually a critical point is reached and both the sample population and the average across all 50 competing populations goes to 100% cooperation.

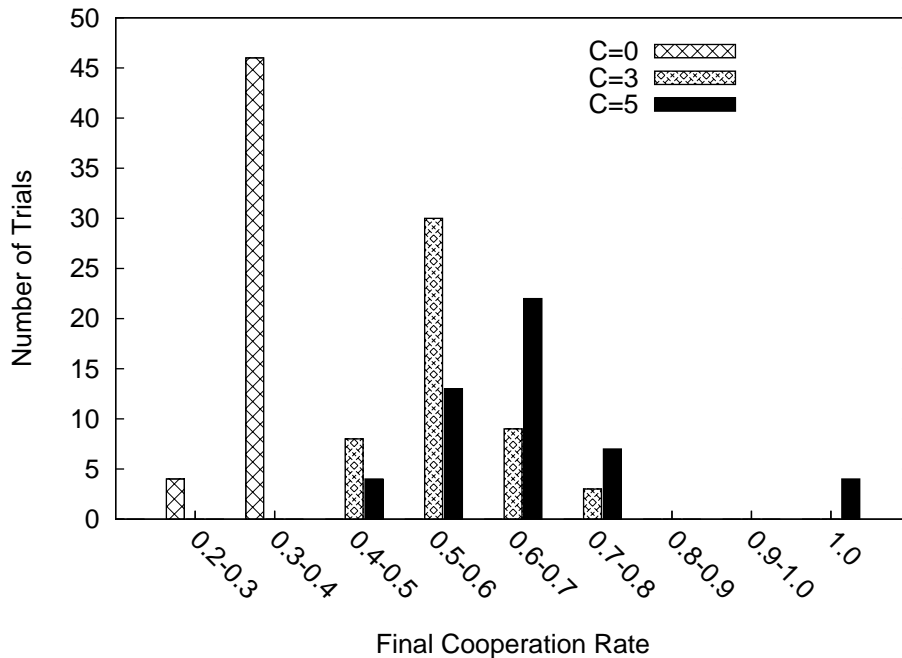


Figure 7: Distribution of average cooperation rates in the final iteration with a neighborhood of 5 and the three rates of inter-population competition ( $C = 0, 3,$  and  $5$ ). With no inter-population competition ( $C = 0$ ) final cooperation rates are relatively low ( $0.2-0.4$ ). As inter-population rates increase ( $C = 3$ ) final cooperation rates increase ( $0.4-0.8$ ) and eventually (for  $C = 5$ ) some populations reach the stable state of full cooperation.

and the average cooperation rate across all 50 populations roughly follows those fluctuations. The figures show that high levels of cooperation often extend to the whole set of populations and persist for many iterations, even before a stable state of nearly 100% cooperation is reached. E.g. in Figure ?? there is a period of relatively high average cooperation around iteration 40,000 that lasts for several thousand iterations.

Spector and Klein have explained the mechanism by which such high levels of cooperation arise [4, 11]. They suggest that in a population mostly consisting of non-sharing individuals a single mutant with a relatively high tolerance may appear and donate to one or more of its neighbors. With their above average resources (due to the donation) these neighbors reproduce and their offspring populate the immediate neighborhood. Because these offspring have identical (or at least similar) tags they form a sharing nucleus that quickly spreads throughout the population. Of course, defectors arise almost as quickly and the cooperation rate drops back to (near) zero. Spector and Klein also showed that with smaller neighborhoods the similar members in these sharing nuclei




Figure 8: Distribution of average cooperation rates in the final iteration with a neighborhood of 10 and the three rate of population level competition ( $C = 0, 3,$  and  $5$ ). Compared to Figure 5 cooperation rates are generally lower, e.g. results for  $C = 0$  and  $C = 3$ , but the steady state of 100% cooperation is reached more frequently, occasionally for  $C = 3$  and always for  $C = 5$ .

are more likely to be chosen to interact, increasing their chance of donating to each other and spreading throughout the population. Thus, with no second level sharing smaller neighborhoods lead to higher average cooperation rates (as seen in Figure 2).

The difference between the trials with  $C = 0, 3,$  and  $5$  is in how frequently, and how sustained, high levels of cooperation are. For  $C = 0$  the sample population occasionally exhibits very high levels of cooperation, but they are very brief and don't significantly effect the average cooperation rate. For  $C = 3$  (three population level competitions per iteration) periods of high cooperation are more frequent, and occasionally spread to a reasonable percentage of the competing populations, occasionally producing high average (across all 50 populations) cooperation rates - e.g. around iteration 40,000. This occurs because the competition between populations creates a mechanism for high levels of cooperation to spread between populations.

For  $C = 5$  the situation is magnified and is reversed from the  $C = 3$  case. Now high average cooperation is the norm and periods of low average cooperation are rare. In Figure ?? one such period of low average cooperation appears

around iteration 27,000. If high enough cooperation rates spread to enough of the competing populations the whole system moves into a state of near 100% cooperation. This also occurs for  $C = 3$ , but less frequently.

To determine how often and under what circumstances the stable state of full cooperation is reached, Figures 5 and 6 present the distribution of average cooperation rates in the final (440,000th) iteration. These figures confirm that the smaller neighborhoods lead to higher average levels of cooperation. However, as competition between the populations plays an increasing role the populations with a larger neighborhood are more likely to reach the critical value of 80% cooperation and move from there into the stable state of pure cooperation.

It is also clear that the populations that have reached the stable state distort the averages seen in Figure 2. E.g. the relatively high average cooperation rate with a neighborhood of 10 and  $C = 4$  observed in Figure 2 is because nearly 50% of the trials have reached the stable state. In those trials that haven't reached a stable state it is clear a neighborhood of 5 produces a higher cooperation rate.

### 4.3 Tagless Results

To verify the importance of the tags, even in the presence of inter-population competition, we also ran trials without tags. In this final set of experiments an individual's decision to donate is based only on their own tolerance value: individuals with a tolerance greater than 0.5 always donate (regardless of the relative values of their and the receivers tags) and individuals with a tolerance less than 0.5 never donate. Thus, cooperation can still evolve. For these trials the neighborhood size is 10 and the inter-population competition rate is 5. These are the values that lead to the highest cooperation rates in our previous experiments.

Figure 7 shows the average cooperation rate and the average tolerance rate when tags are ignored. Initially the cooperation rate is roughly 0.5%, as expected in populations consisting of individuals with tolerances randomly distributed between 0 and 1.0. However, the cooperation rate quickly drops to 0 and stays there (data out to 4000 iterations show no change in the cooperation rate). Figure 7 also shows that the average tolerance drops to just under 0.25. Only individuals with a tolerance above 0.5 donate, so this confirms that the populations are evolving to all defectors, even in though there is a high level on inter-population competition.

## 5 Discussion

Recently there has been considerable interest in how cooperation evolves. We have addressed some of the fundamental questions regarding the evolution of cooperation in an environment with multiple levels of competition. Our results lead to several significant conclusions.

First, it is clear that in this model cooperation rates typically undergo rapid and extreme changes. Looking at the average cooperation rate within a popula-




Figure 9: Cooperation rate (left) and average tolerance (right) when tags are ignored. Individuals always donate if their tolerance is greater than 0.5 and otherwise never donate. Without the ability to discriminate based on relative tag values cooperation isn't sustained.

tion or across multiple populations can be misleading because it obscures these rapid fluctuations in the cooperation rate. In particular, the changes in average cooperation rate observed by Spector and Klein as a function of the neighborhood size and the mutation rate are a result of changes in the frequency and persistence of the high cooperation periods, not the result of changing a relatively stable cooperation rate. This result suggests the need for a change in how cooperation is viewed and measured in, at least some, models. If cooperation levels are widely fluctuating, as in these experiments, measuring cooperation in terms of the frequency with which it evolves in the population and the stability of cooperation once it evolves may be much more informative than simply measuring the average rate of cooperation.

Second, our results confirm that multi-level competition can increase the average cooperation rate. In fact, multi-level competition has two significant affects on the average cooperation level.

- Multi-level competition can increase the average cooperation rate by increasing the frequency and duration of the periods of high cooperation.
- Multi-level competition can produce a stable state in which all of the competing populations have near 100% cooperation. Once this state is reached defection becomes an unsuccessful strategy.

Because of the nature of the second level competition and replacement being used (in which a population with fewer resources is completely replaced by a population with more resources) the results are most applicable to cases where cooperation can be learned. E.g. where one group can observe that in another group cooperation is beneficial and can copy the cooperation strategy. Considerable additional research is needed to determine the role of multi-level competition on the evolution of cooperation when other models of inter-group competition are used.

## References

- [1] Elizabeth Pennisi. How did cooperative behavior evolve? *Science*, 309(5731):93, 2005.
- [2] Martin A. Nowak. Five rules for the evolution of cooperation. *Science*, 314(5805):1560–1563, 2006.
- [3] Rick L.Riolo, Michael D. Cohen, and Robert Axelrod. Evolution of cooperation without reciprocity. *Nature*, 414(6862):441–443, 2001.
- [4] Lee Spector and Jon Klein. Genetic stability and territorial structure facilitate the evolution of tag-mediated altruism. *Artificial Life*, 12(4):553–560, 2006.
- [5] John Holland. *Hidden Order: How Adaptation Builds Complexity*. Addison Wesley, 1995.
- [6] David Sloan Wilson. Weak altruism, strong group selection. *Oikos*, 59(1):135–140, 1990.
- [7] Jeffrey A. Fletcher and Martin Zwick. Strong altruism can evolve in randomly formed groups. *Journal of Theoretical Biology*, 228:303–313, 2004.
- [8] John w. Pepper. Relatedness in trait group models of social evolution. *Journal of Theoretical Biology*, 206:355–368, 2000.
- [9] Gilbert Roberts and Thomas N. Sherratt. Does similarity breed cooperation. *Nature*, 418:499–500, 2002.
- [10] Rick L.Riolo, Michael D. Cohen, and Robert Axelrod. Does similarity breed cooperation. *Nature*, 418:500, 2002.
- [11] Lee Spector and Jon Klein. Multidimensional tags, cooperative populations, and genetic programming. In Terence Soule Rick Riolo and Bill Worzel, editors, *Genetic Programming Theory and Practice IV*, pages 97–112. Springer US, 2007.
- [12] Timothy Killingback, Jonas Bieri, and Thomas Flatt. Evolution in group-structured populations can resolve the tragedy of the commons. *Proceedings of the Royal Society B*, 273:1477–1481, 2006.
- [13] Samuel Bowles, Jung-Kyoo Choic, and Astrid Hopfensitzd. The co-evolution of individual behaviors and social institutions. *Journal of Theoretical Biology*, 223:135–147, 2003.
- [14] Vincent A. A. Jansen and Minus van Baalen. Altruism through beard chromodynamics. *Nature*, 440(30):663–666, 2006.