# Institutions and Cooperation in an Ecology of Games

# Paul E. Smaldino\*,\*\*

Johns Hopkins University

## Mark Lubell<sup>†</sup>

University of California, Davis

Abstract Social dilemmas have long been studied formally as cooperation games that pit individual gains against those of the group. In the real world, individuals face an ecology of games where they play many such games simultaneously, often with overlapping co-players. Here, we study an agent-based model of an ecology of public goods games and compare the effectiveness of two institutional mechanisms for promoting cooperation: a simple institution of limited group size (capacity constraints) and a reputational institution based on observed behavior. Reputation is shown to allow much higher relative payoffs for cooperators than do capacity constraints, but only if (1) the rate of reputational information flow is fast enough relative to the rate of social mobility, and (2) cooperators are relatively common in the population. When these conditions are not met, capacity constraints are more effective at protecting the interests of cooperators. Because of the simplicity of the limited-group-size rule, capacity constraints can also generate social organization, which promotes cooperation much more quickly than can reputation. Our results are discussed in terms of both normative prescriptions and evolutionary theory regarding institutions that regulate cooperation. More broadly, the ecology-of-games approach developed here provides an adaptable modeling framework for studying a wide variety of problems in the social sciences.

#### Keywords

Public goods, social dilemma, group size, reputation, capacity constraints, agent-based

A version of this paper with color figures is available online at http://dx.doi.org/10.1162/ artl\_a\_00126. Subscription required.

## I Introduction

This article develops a computational model to analyze how different institutional arrangements affect cooperation in an ecology of public goods games. The analysis is motivated by the fact that in most real-world settings, political and economic outcomes emerge from actors engaging in multiple strategic interactions in different types of venues. How cooperation is maintained in such complex networks of games and actors is therefore a pressing concern. However, most of the vast literature on cooperation and other types of strategic interaction focuses on situations where actors play one game at a time, or within one unstructured population [5, 6, 8, 11–13, 17, 19, 22, 23, 26]. The model developed here begins to build a more general theoretical framework that extends the original idea of an ecology of games developed by the sociologist Norbert Long [28], who applied it to local political economics.

<sup>\*</sup> Contact author.

<sup>\*\*</sup> Center for Advanced Modeling in the Social, Behavioral, and Health Sciences and Department of Emergency Medicine, Johns Hopkins University, Baltimore, MD 21209. E-mail: paul.smaldino@gmail.com

<sup>†</sup> Department of Environmental Science and Policy, University of California, Davis, CA 95616. E-mail: mnlubell@ucdavis.edu

Our model has both theoretical and substantive motivations. On the theoretical side, it is important to compare the processes affecting cooperation in an ecology of games with those in other types of cooperation models. A variety of spatial and network models have been presented with an aim of explaining cooperation in structured populations (reviewed in [37, 42, 48]). The unifying feature of these models is that cooperators outperform defectors when cooperators *positively assort*—that is, when cooperators are more likely to interact with one another than expected by their representation in the population at large [19]. Models are therefore distinguished by the mechanisms that generate assortment. Here, we will investigate institutional mechanisms for generating assortment in an ecology of games. We will discuss these in the context of previous modeling work, while also differentiating the ecology-of-games framework from other modeling approaches.

On the substantive side, examples of multiple games are ubiquitous from the scale of daily life to the most fundamental economic and political processes in society. For example, Jimmy may be involved in work, school, church, clubs, and political organizations. Jimmy may go to both school and church with Sally, who along with Rita does charity work and sits on the city council. In the context of environmental governance, a recent analysis of the San Francisco Bay identified over 380 actors and 116 different water policy games, in which the average policy actor participated in three or more games [29]. In the context of economic activity, the ecology-of-games idea is implicit in network analyses of interlocking corporate directorates [22, 24] and innovation networks where different companies engage in dynamic partnerships over time [18, 30]. But these examples are just the tip of the iceberg—recognizing the existence of an ecology of games usually requires only a brief moment's reflection in most research contexts.

Our analysis advances these theoretical and substantive issues in a number of ways. First, we develop a modeling framework that is capable of analyzing multiple games. Second, we show that reputation plays a similar role in the ecology of games setting to the one it plays in the more traditional repeated game setting [25, 36], based primarily on the mechanism of positive assortment. Third, and perhaps most importantly, we outline key issues involved with comparative institutional analysis, where the question becomes *how* different types of institutions promote cooperation given different environmental and social contexts. Two institutions that both facilitate the positive assortment of cooperators may differ in their resultant degrees of positive assortment, in the speed with which they can generate assortment, and in their effectiveness in different demographic pools (e.g., the frequency with which individuals tend to be cooperators). We investigate these factors and identify conditions in which one institution or another best promotes cooperation in an ecology of games.

# 2 Cooperation in an Ecology of Games

As in previous work [45], we modeled the ecology of games as groups of agents adopting pure strategies of either defect or cooperate in multiple public goods games, who leave games where they experience losses and attempt to join games with higher payoffs. The public goods game is a well-known model for social dilemmas that pits group benefit against individual greed. Each time the game is played, an individual has resources c = 1, which he or she may either contribute to a group pot (*cooperate*) or keep for himself or herself (*defect*). The total amount in the pot is then increased by some factor and redistributed among all players. Groups therefore do best when everyone cooperates, but individuals maximize their payoffs by defecting.

Here, we are specifically interested in simple institutions that allow cooperators (altruistic individuals) to outcompete defectors (free riders) in an ecology of public goods games. By institution, we mean formal and informal rules that govern human interaction [34]. We assume that individuals playing a game may employ certain rules to determine whether or not a given applicant will be allowed to join their game. Whether cooperation is promoted in the ecology of games therefore depends on the effectiveness of these institutions in permitting cooperators to outperform defectors, which in turn is a function of positive assortment among cooperators. In this article, we investigate two such institutions: *capacity constraints* and *reputational exclusion*. Capacity constraints impose a fixed upper limit on the number of agents that can join any particular game. Capacity constraints were investigated previously by Smaldino and Lubell [45], and were shown to be sufficient to allow cooperators to assort to the extent where they outperform defectors. This is due to cooperators leaving games in which they are heavily exploited by defectors; capacity constraints then prevent defectors (as well as further cooperators) from joining games with high levels of cooperation. This finding fits with other theoretical work showing that limited patch size promotes cooperation in evolutionary contexts [41, 46], and demonstrates how institutions that promote inclusiveness without an explicit mechanism for positive assortment of cooperators risk being overrun by defectors.

We compare capacity constraints with an institution of reputation-based exclusion that allows current players in a game to exclude or admit new participants on the basis of a known behavioral history. Reputation institutions are common in many realms of social life, and allow selective admission of cooperators to a game without increasing the number of defectors. Reputation provides an alternative to capacity constraints for sorting cooperators and defectors, without limiting the potential total number of actors interacting in a group. As we will show, the effectiveness of reputation depends on the reliability of information about past behaviors, and how quickly players move from one game to the next.

A simple way to operationalize reputation is known as *image scoring* [31, 35, 36, 47, 52]. An individual's image score increases when he or she is observed defecting. We assumed that image scores are subjective—that is, individual A's image score increases in B's eyes when B observes A cooperating, but does not effect C's image score for A if C is not present to witness the behavior. Theoretical models have shown that image scoring may be of limited robustness as a strategy for enforcing cooperation when an agent uses a co-player's image score to decide game play behavior (cooperate or defect) in the presence of errors [40]. Here, image scoring is used to guide assortment rather than strategy choices, and as discussed below is highly effective at promoting cooperation in the presence of perceptual errors over a wide range of conditions. Experimental studies of single-game dynamics have shown that humans often use image scoring as an effective means of enforcing cooperation [7, 10, 32, 47, 49].

Individuals in our agent-based model could play multiple games simultaneously, and the processes of leaving current games and joining new ones were decoupled, in contrast to many alternative models of games on dynamic networks [20, 21, 39, 42-44, 51] (but see also [9] for an exception). By decoupling leaving and joining, and by making games distinct entities rather than their being defined by a given individual's social neighborhood, we can take account of the fact that individuals interact in multiple settings with overlapping co-players without being restricted to a fixed lattice or network structure. Additional differences between the basic ecology-of-games model and previous models of cooperation are discussed in [45]. Our model of reputational exclusion also differs from previous models of reputation and cooperation in two important ways. First, some models implement a strategy of contingent movement away from current games based upon the previously observed actions of co-players [1, 2, 16, 44] or the payoffs resulting from those actions [3], with movement toward new games being randomly directed. In our model, movement choices are based solely on the perplayer output of potential games and the net profit from current games, although the emergent behavior is quite similar to active movement toward cooperators and away from defectors. Second, players in some models adjust their game behavior according to their own reputations and/or those of their co-players [14, 15, 27, 31, 35, 40]. Players in our model play pure strategies, and we consider only the sociostructural reorganization within individual lifetimes. Image scores are used by agents to decide whether to allow or forbid a target individual to participate in additional games.

Reputation allows game players to selectively bar individuals with poor reputations from joining their games, while an individual in good standing may still gain entry into games in which both he or she and his or her new co-players stand to profit by his or her inclusion. In many formal settings, it is often far more difficult to terminate a current participant's membership than to prevent that individual from joining in the first place, and so we did not include a mechanism for ousting an unwanted player from a game. North [34] distinguishes between formal institutions, which include laws, regulations, and policies, and informal institutions, which are based on social conventions and behavioral norms. The reputationbased institution we consider here is easily seen as a hybrid: a formal institution (voting) backed by an informal institution (keeping track of others' behaviors and preferring cooperators). Although we primarily view capacity constraints as a formal institution (because it would be a straightforward rule to implement), a hard limitation on group size could also function as an informal institution in the form of a "fast and frugal" heuristic [21] for group decision making: When assessing whether to admit an unknown applicant, admit if and only if the number of current players is below some maximum.

#### **3 Model Description**

A population of N agents was initially distributed among M public goods games. Agents played pure strategies of cooperate or defect, with f representing the frequency of cooperators, and were initially placed randomly in  $\gamma_0 M$  games,  $0 < \gamma_0 \le 1$ . Each agent i also had a subjective image score for every other agent j,  $x_{ij}$ , which was initially zero for all agents. If i observed j to cooperate (i.e., i and j were co-players in at least one game),  $x_{ij}$  increased by one. If i observed j to defect,  $x_{ij}$ decreased by one. Agents' attention was set narrowly; each co-player was observed only once per time step, so there was no gain in social information from sharing more games with a given co-player. Each subjective image score could therefore only change by one unit each time step, regardless of how many games two agents shared.

Some past models of reputation have assumed that precise information about each individual's past behavior is shared among everyone in the population, resulting in a public image score or social standing for each individual [16, 27, 35, 36, 47]. Here, individuals only have private information on those whom they have directly observed, and information is socially aggregated within a given game through the pooling of inclinations (described below), which represents a form of information condensation.

Each time step, all active public goods games were played. The per-agent payout for each game G was

$$\pi_G = \frac{r}{n_G} \sum_{i=1}^{n_G} \frac{s_i}{m_i},$$
(1)

where *r* is the game production,  $n_G$  is the number of players in game *G*,  $m_i$  is the number of games being played by agent *i*, and  $s_i = c$  if *i* was a cooperator and zero if *i* was a defector. For example, a cooperator playing four games will contribute c/4 to each game, while defectors contribute zero. The total contributions to each game are then multiplied by *r* and distributed among all players, so that defectors enjoy the benefits of contributions from cooperators without contributing themselves. Note that if *r* is the same for all games (which was assumed), the system would reduce to a single public goods game if all agents were to play all games simultaneously, although this never occurred when agents employed either capacity constraints or reputational exclusion.

After payoffs were obtained, each agent attempted to join a new game with probability  $\gamma$ , where  $0 < \gamma \leq 1$  determines the rate of agent *social mobility*. This is the likelihood that an individual attempts to leave and join new games relative to the time of meaningful game interactions (time steps). Individuals could only join games that were being played by a current co-player, representing the fact that information about the existence of social opportunities spreads through networks. Relaxing this last constraint did not qualitatively affect our results.

When joining was attempted, an agent considered all games with positive payouts ( $\pi_G > 0$ ) and attempted to join one with probability proportional to its relative payout. The current players of this

game then employed one of two institutions to decide collectively whether to allow the applicant agent to join their game.

## 3.1 Capacity Constraints

Using this simple institution, all games have a maximum number of players,  $n^* < N$ . A new applicant was allowed to join a game if and only if the current number of players was less than  $n^*$ .

#### 3.2 Reputational Exclusion

Here the current players of a game collectively decided whether to allow the applicant agent j to join their game. Each current player i had an *inclination*  $I_{ij}$  toward accepting j, which was based on i's subjective image score of i and given by the logistic function

$$I_{ij} = \frac{1}{1 + e^{-\alpha(x_{ij} + \varepsilon)}},\tag{2}$$

where  $0 \le \alpha \le 1$  was the rate of *reputation convergence* and  $\varepsilon$  was an Gaussian noise term with a mean of zero and a standard deviation of one. The term  $\varepsilon$  represents potential errors in observation as well as variation in decision certainty due to random events. Applicant agent *j* was admitted to game *G* with probability equal to the average inclination of all *G*'s current players toward admitting *j*. Reputation convergence represents the speed at which information about observed acts of cooperation or defection decreases the uncertainty in decisions about group admittance (and therefore the inclination to admit an applicant or not), and thus influences the probability that positive image scores translate to an inclination to admit close to one and negative image scores translate to an inclination close to zero (Figure 1). Large  $\alpha$  meant that only a few observations were necessary; small  $\alpha$  meant that many observations were required. The rate of reputation convergence can be interpreted as an indicator of the time it takes to gather reliable and predictive information about an individual for social decision making, or as the maximum rate at which reliable information can be acquired.

With probability  $\gamma$ , each agent then entertained the possibility of leaving one of its current games, a decision that was independent of whether the agent had just attempted to join a new game. In this case, an agent considered all of its current games, excluding any just joined, and left the one with the

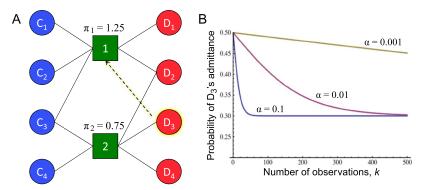


Figure I. (A) Schematic diagram of example ecology with eight agents and two games with r = 2.5. Edges represent membership. Agent D<sub>3</sub> attempts to join game I, in which only agents D<sub>2</sub> and C<sub>3</sub> have observed D<sub>3</sub>'s behavior. (B) Reputational exclusion. Assuming k rounds of play prior to this joining attempt and no relevant prior movement, the probability of D<sub>3</sub>'s being admitted to game I, based on Equation 2 and ignoring error, is given by

$$Pr(yes \text{ to } D_3) = \frac{I}{5} \left[ \frac{3}{2} + 2 \cdot \frac{I}{I + e^{-\alpha(-I \cdot k)}} \right].$$

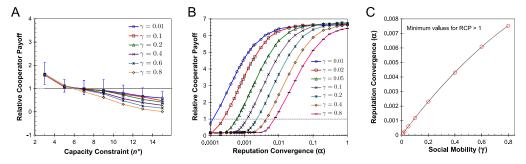


Figure 2. (A) RCP as a function of the capacity constraint  $n^*$  for several values of social mobility  $\gamma$ . Error bars are the standard deviation for the  $\gamma = 0.01$  runs. (B) RCP as a function of reputation convergence  $\alpha$  for different values of  $\gamma$ . The more quickly and reliably reputation information can be used, the better cooperators do. The more socially mobile individuals are, the more important it is for reliable reputation information to be able to amass rapidly. This relationship is very close to linear, as noted for the paired values of  $\alpha$  and  $\gamma$  for when RCP > 1 (C).

lowest payout if and only if that payout was less than or equal to the individual's contribution to that game. In other words, individuals only left games that did not yield positive returns. We assumed that the incentives for joining or leaving games depended only on each game's independent payoff; for simplicity, we ignored other factors that could potentially influence real-life group choice decisions, such as a group's historical reputation and the influence of specific members beyond their tendency to cooperate.

The process of sociostructural reorganization was continued until an equilibrium was reached. We considered the system to have reached equilibrium if either (a) there was no variation in payoffs among cooperators *and* no variation in payoffs among defectors, or (b) there were no changes in the average cooperator payoff, average number of agents per game, or average number of games per agent for 100 time steps. Our results are averaged from 100 runs of each condition. Unless otherwise stated, we used values of r = 2.5, N = 100, M = 100, and  $\gamma_0 = 0.1$ . As in previous work [45], our primary metric for assessing the success of cooperators is the global relative cooperator payoff (RCP). This is simply the average payoff to cooperators divided by the average payoff to defectors, so cooperators do better than defectors when RCP > 1.

#### 4 Comparative Institutional Analysis

Reputation allowed cooperators to assort while keeping out defectors, thereby allowing much higher relative payoffs for cooperators than either an unconstrained scenario without reputation, or one with capacity constraints alone. For example, with N = M = 100 and r = 2.5, the best RCP that could be expected under capacity constraints was about 1.6, meaning that cooperators performed about 60% better on average than defectors. With reputation, the maximum relative cooperator payoff we found was approximately RCP = 6.7, a more than fourfold advantage over capacity constraints (Figure 2A, B). The ability of reputation to produce favorable conditions for cooperators was dependent on the rate of reputation convergence being fast enough for the given level of social mobility in the population. If agents had more opportunities to leave and join games relative to the time scale of game payoffs, then reputation needed to be more efficient (i.e., fewer observations were needed for certainty of decisions) in order to ensure cooperation (Figure 2C). For example, for social mobility of  $\gamma = 0.6$ , reputation had to converge at a rate of at least  $\alpha = 0.0061$  in order for cooperators to outperform defectors. Capacity constraints were relatively insensitive to social mobility, since they relied primarily on agents leaving games without positive net payoffs.

Although reputation could produce a larger cooperator advantage than capacity constraints, capacity constraints enabled the system to settle into an equilibrium much more quickly than reputation. For runs in which cooperators had an eventual advantage (i.e., RCP > 1), the variation

in time to equilibrium—in which agents no longer joined new games or left old ones—was primarily a function of social mobility ( $\gamma$ ), rather than the strictness of capacity constraints ( $n^*$ ) or the efficiency of reputation ( $\alpha$ ). Reputation, however, took longer to reach equilibrium than capacity constraints by an order of magnitude (Figure 3) for all levels of social mobility. Some of this time was taken up by cooperators slowly and stochastically joining games until there were no more profitable games left to join, well after reputation had converged to prohibit defectors from joining most desirable games. When we considered only the time necessary for cooperators to have a marginal average advantage over defectors (RCP > 1), we saw that reputation could achieve this much more rapidly, but still not nearly as quickly as capacity constraints.

Because agents play pure strategies without evolution in our model, the global frequency of cooperators, f, is a free parameter. Our model predicts that reputation will outperform capacity constraints as an institution for promoting cooperator success when cooperators are common, but not when they are rare (Figure 4). Under reputation, RCP increased monotonically with f, as this simply meant fewer defectors and more cooperators in each game (Figure 4A). When f was low, the stochastic process meant more chances for defectors to trickle into games. Early on, while admission decisions involved a lot of uncertainty, defectors joined many games simply because most applicants were defectors, and the games most likely to be joined were those containing the most cooperators. In general, when cooperators were rare, reputation could not promote cooperation unless  $\alpha$  was very large. Capacity constraints, on the other hand, produced maximal RCP when cooperators were in the slight minority (Figure 4B). The effect was due to the fact that the number of players per game was limited, and players thereby became constrained in the number of games they could join due to lack of available spaces. This restricted the number of games in which a cooperator could play. When cooperators were common, defectors tended to be surrounded by cooperators, and so their relative payoffs were quite high. Conversely, when cooperators were rare, capacity constraints could still promote cooperation, because most defectors were excluded from games that had filled up with cooperators.

#### 5 Reputation in More Detail

#### 5.1 Example Dynamics

Precisely how capacity constraints promote cooperation is explored in detail in [45]. Here we detail the model dynamics when agents control game entry through reputational exclusion. Figure 5 shows

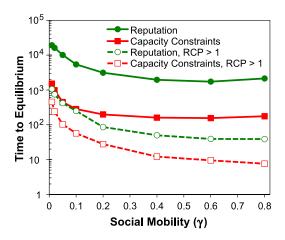


Figure 3. Time to equilibrium (solid lines) averaged across all values of  $\alpha$  (for reputation) and  $n^*$  (for capacity constraints) for parameter conditions where more than 75% of runs ended with RCP > 1. The system took far longer to converge to an equilibrium under reputation than under capacity constraints. However, if we consider only the time to reach RCP > 1 (dashed lines), then the difference between reputation and capacity constraints is smaller.

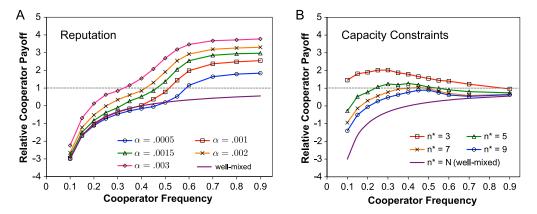


Figure 4. Relative cooperator payoff at equilibrium as a function of the global cooperator frequency f for reputation (A) and capacity constraints (B). RCP increased monotonically with f, because cooperators could assort without limit and still be relatively free of defectors. Note that for low-middle frequencies of cooperators, capacity constraints outperform reputation. For all runs shown,  $\gamma = 0.1$ .

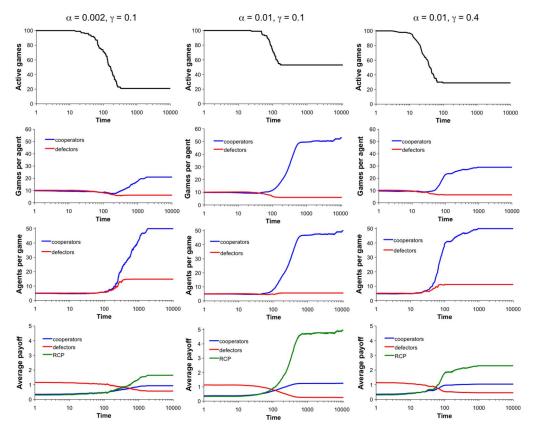


Figure 5. Example game dynamics. From top to bottom, the graphs show the number of active games (games with at least one player), the average number of games played by each agent type, the average number of cooperators and defectors per active game, and the average payoffs for cooperators and defectors. The middle column can be compared with the columns on the left and right to see the effects of decreased reputation convergence or increased social mobility, respectively.

example dynamics for three runs of the model with reputation. Before reputation converged, there was not enough observation-based information for agents to effectively identify defectors and bar them from admittance; agents were admitted to games largely at random and left games if there were too few cooperators. The number of active games dropped significantly before reputation had an effect, as agents left games that did not yield positive payoffs (top row). Most of the games remaining were those in which there were enough cooperators to make the public goods payoff above the average cost of contribution, although defectors may still have had the overall advantage, since they did not incur costs. Once reputation converged, cooperators continued to leave games in which they did not do well, and joined games in which average payoffs were high (and which tended to have more cooperators present). This led to more games per cooperator and fewer games for the average defector (second row). The third row shows that there were far more cooperators than defectors in the average game. The plateau in defectors per game marks the point at which no further defectors were permitted to join games, although the number of cooperators in games continued to climb. These factors caused cooperator payoffs to rise and defector payoffs to fall (bottom row). The figure also shows how the success of cooperators decreased when either the rate of reputation convergence decreased (left column) or social mobility increased (right column).

# 5.2 Excess Assortment

Reputation allowed a game's current players to admit cooperators and keep out defectors. In any given game, however, defectors received larger payoffs than cooperators, since all agents received the same return but defectors contributed nothing. Cooperators' ability to outperform defectors lay therefore not merely in the fact that a game had more cooperators than defectors, since defectors do better in any one such game. Cooperators' success lay in their ability to join many more games than defectors, and therefore receive larger cumulative payoffs across games. The ecology-of-games perspective highlights the importance of analyzing the benefits of cooperation at the scale of multiple games, rather than the scale of individual games where defectors always win.

Because each game's current players had a specific mechanism to allow cooperators into games and keep defectors out, cooperators were positively assorted, and defectors were negatively assorted (i.e., a given co-player was more likely to be a cooperator). As in previous work [41, 45], we can quantify excess assortment as the probability above chance at which an agent encounters another agent of the same kind. Cooperators were spread among more games than defectors-indeed, they eventually joined most active games-and so cooperators' tendency to positively assort was equaled (or nearly equaled) by defectors' tendency to negatively assort. In other words, the probability that two randomly chosen agents in a given game were both cooperators exceeded the global cooperator frequency to the same degree that the probability that the two agents were both defectors was exceeded by the global defector frequency. The total excess assortment across all agents was therefore negligible. We therefore looked at the excess assortment for cooperators only, which we computed as the average per-game cooperator frequency at equilibrium minus the global cooperator frequency. Figure 6A shows that the advantage to cooperators was driven entirely by their ability to positively assort, which was in turn driven by the ability of reputation to keep defectors out of games. The faster reputation converged relative to the population's social mobility, the more effectively defectors were kept out of games, and therefore the fewer defectors were found in active games at equilibrium (Figure 6B). The assortment result occurs primarily through the barring of defector entry, since cooperators tended to join all active games if possible, though this was not always completely possible for low levels of social mobility, since agents' social networks might not grow large enough to make every game available to them (Figure 6C).

## 5.3 Number of Active Games

The more effective reputation was at promoting cooperation, the more quickly defectors could be kept from joining additional games. This led to more active games remaining at equilibrium (Figure 7). Because defectors always remained in a game as long as there was at least one cooperator (as it cost

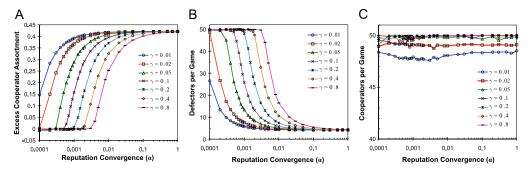


Figure 6. (A) Excess cooperator assortment, and (B) the average numbers of defectors and (C) cooperators per game at equilibrium with reputation, as a function of  $\alpha$  for several values of social mobility  $\gamma$ .

them nothing and yielded a positive return), the driving factor in a game becoming inactive was the exodus of cooperators. This occurred when the ratio of cooperators to defectors was too low, with the precise threshold depending on both the game production (*r*) and the number of games among which each cooperator divided his or her resources. Games that initially had too few cooperators could become inactive unless more cooperators joined before the extant cooperators left. Another way that games became inactive was by an influx of defectors, which drove the payoff of cooperators down. Once too many defectors joined, the game became a losing investment for cooperators, and they started to leave. A slower reputation convergence rate allowed more defectors to join games, which meant more games became inactive as cooperators fled. This compounded the loss for cooperators, as they were able to receive payoffs from fewer total games. An analogous result was seen with capacity constraints in the absence of reputation: When agents were allowed to join games, more games became inactive [45].

#### 5.4 Game Production

The specific game dynamics were dependent not only on  $\alpha$  and  $\gamma$ , but also on the game production *r*. Game production dictated the quantitative return on game investments, as well as the ratio of cooperators to defectors for which cooperators continued to receive positive net payoffs, though as mentioned this ratio also depended on the number of games among which each cooperator divided his or her contributions. The most important result is that, with reputation, the relative cooperator payoff

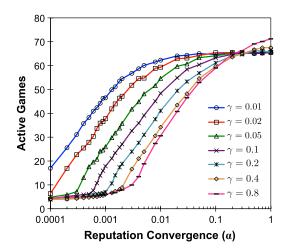


Figure 7. The average number of active games at equilibrium as a function of reputation convergence, for several values of social mobility.

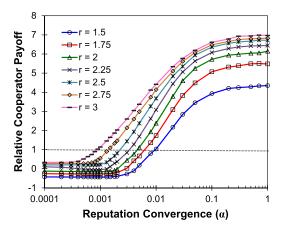


Figure 8. Varying the production function r. For higher values of r, RCP increased across the board. Note that the differences between runs tended to be greatest for large values of  $\alpha$ . For all runs here  $\gamma = 0.2$ .

increased monotonically with the game production (Figure 8). This contrasts with the model under capacity constraints and no reputation [45]. In that case, the tight restriction of game size meant that an increase in r could decrease the minimum number of cooperators per sustained game, so that cooperators would stay in games they would have otherwise left (if it had smaller r), which actually decreased their relative payoff. The ability of cooperators to join games without limit is therefore an appealing factor of reputation, as it does not penalize cooperators for participating in more productive games.

#### 5.5 Limited Abilities to Play Multiple Games

Throughout our analysis, we have assumed that the only restriction to an agent joining a new game is whether or not joining is permitted by the game's current players. The ability to participate simultaneously contrasts starkly with most prior analyses, which assume that individuals participate in only one game at a time. However, it is not necessarily a realistic assumption that agents would pay no cost to participate in multiple games. Social interactions often require costs in the form of time, money, and maintaining social relationships. In this case, for example, an agent might decide to leave a game with a low but positive net payoff because participation in such a game incurred an opportunity cost in the lost ability to participate in potentially higher-paying games.

One of our goals in making our model was to give agents only limited cognitive decision-making capabilities in order to demonstrate how cooperation in an ecology of games emerges from simple

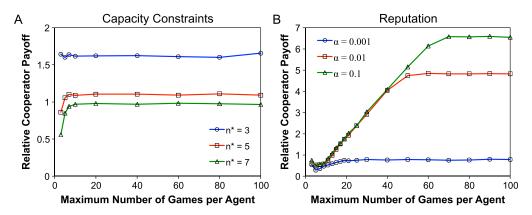


Figure 9. Limitations to the number of games an agent could play, for both capacity constraints (A) and reputational exclusion (B). For all runs here  $\gamma = 0.1$ .

but reasonable behavior. Nevertheless, it was important to recognize more precisely the reasons why individuals might need to adopt cognitive and institutional mechanisms for handling opportunity costs related to limited abilities to play in multiple games. We implemented a simple limitation on the maximum number of games an agent could play. This limitation was termed "budget constraints" in [45], and shown to be ineffective at promoting cooperation in the absence of other institutions. In this extension, the model is as described above, except that joining a new game was not attempted if an agent was already playing the maximum number of games per agent. This limitation had only a very small influence on the effects of capacity constraints (Figure 9A). When  $n^* = 3$ , there was no statistically significant effect at all. For larger game capacities, there were effects only when agents were restricted to playing in very few games. This lack of effect was due to the fact that, under capacity constraints, even with no limitations to the number of games per agent, the average number of games per agent was consistently less than  $n^*$ . In contrast, limiting the number of games per agent severely decreased the potential advantage to cooperators when reputational exclusion was the active institution (Figure 9B). This was because the ability of reputational exclusion to promote cooperation lies partly in its ability to allow cooperators to join many games. Reputational exclusion was still effective at promoting cooperation when there were limits to the number of games each agent could join, but it was less so than in the absence of that restriction.

# 6 Discussion

Which institutions are most effective at promoting cooperation will depend on a variety of conditions. These conditions include the rate at which individuals move between social groups, the rate at which reliable social information can be acquired, and the overall frequency of cooperative individuals in the population. We have shown reputational exclusion to be a potentially excellent mechanism for promoting cooperation in an ecology of games. Not only does it allow cooperators to outperform defectors, it can achieve much higher relative cooperator payoffs than can capacity constraints, because cooperators can assort while keeping out defectors. Reputation is only effective, however, if the rate at which observations can usefully guide meaningful decisions (i.e., reputation convergence) is appropriately fast given the level of social mobility in the population. Furthermore, reputation may only outperform capacity constraints if the frequency of cooperators are rare, instituting capacity constraints is a more effective option for ensuring cooperator assortment. Capacity constraints also require less time than reputational exclusion to generate sufficient assortment for cooperators to trump defectors, regardless of the rate of social mobility. If speed is of the essence, capacity constraints are a more economical solution for promoting cooperation in a dynamic population.

The comparative institutional analysis of the size of the cooperator advantage, and the rate at which it emerges has important consequences for the overall process of institutional change in an environment in which actors play multiple games. Evolutionary theories of institutions [4, 34, 50, 53] suggest that the survival of institutions is proportional to their net benefits. Our models suggest a clear tradeoff between the fast but limited cooperation produced by capacity constraints, and the slower but greater cooperation produced by reputation. If the evolutionary process of institutional selection prioritizes speed and occurs quickly relative to the pace at which cooperation evolves, capacity constraints are more likely to survive. If the process of institutional selection happens at a longer time scale, then reputation institutions are more likely to survive, because there will be enough time for reputation to produce higher relative payoffs for cooperation.

The relative advantage of capacity constraints may be higher when trying to build cooperation from a population with few cooperators from the outset. Our model suggests a potential trajectory for the coevolution of cooperative behavior and cooperation-supporting social institutions. When cooperators are rare, limited group size is an effective way to increase the relative fitness of cooperators, while reputational exclusion is not. As the frequency of cooperators increases and capacity constraints become less effective, the adoption of reputation-based institutions such as reputational exclusion will

facilitate the continued growth of cooperators. How a social system might endogenously experience such institutional change is a crucial question for further models and research.

Our analysis does not consider the overall transaction costs of establishing institutions, which are a key focus of the institutional change literature [34]. Both capacity constraints and reputation require enforcing entry limits of some type, and the reputation institution entails both information costs (assessing reputation) and decision costs (collective decision making). These costs must be subtracted from the benefits of cooperation, and would affect the process of institutional evolution in the ecology of games.

The findings also have normative implications for the overall design of social and political institutions. Democratic political systems express a core value of inclusive participation, but the capacity constraints model requires limiting participation for all actors in order to achieve cooperation. Reputation-based institutions, while still excluding defectors, allow cooperators to achieve higher payoffs but do not constrain the number of games in which cooperators can participate. In addition, the total number of games available for participation is higher with reputation, which provides a rich set of opportunities for social interaction throughout the ecology of games that is reminiscent of de Tocqueville's early observations of American democracy.

We did not formally analyze the use of both capacity constraints and reputation in tandem, as our aim in this article was to lay the basic groundwork for the consideration of each institution separately and to identify their strengths and weaknesses. However, it is easy to imagine scenarios in which hybrids of both institutions might come in handy. For example, capacity constraints alone are effective only when group sizes are restricted to very small numbers. A gradual increase in a game's capacity constraint might be advisable as sufficient time passes for reputational information to accrue.

The ecology of games modeling framework allows for relatively simple models that nevertheless can capture the interplay between individual decision making, group-level institutions, social organization, and mobility. In the future, the framework can be extended to a wider variety of game ecologies. Our findings here display the importance of modeling cooperation in an ecology of games as providing a tool for comparative institutional analysis.

#### Acknowledgments

We thank Michael Makowsky, Timothy Waring, David Broniatowski, and Athena Aktipis for helpful comments. Funding for this work was provided by NSF grant #0921904 and by NIH grant #DP1OD003874.

#### References

- Aktipis, C. A. (2004). Know when to walk away: Contingent movement and the evolution of cooperation. Journal of Theoretical Biology, 231(2), 249–260.
- Aktipis, C. A. (2006). Recognition memory and the evolution of cooperation: How simple strategies succeed in an agent-based world. *Adaptive Behavior*, 14(3), 239–247.
- 3. Aktipis, C. A. (2011). Is cooperation viable in mobile organisms? Simple walk away rule favors the evolution of cooperation in groups. *Evolution and Human Behavior*, *32*, 263–276.
- 4. Alchian, A. A. (1950). Uncertainty, evolution, and economic theory. *Journal of Political Economy*, 58, 211–221.
- 5. Axelrod, R. (1984). The evolution of cooperation. New York: Basic Books.
- 6. Axelrod, R. (1997). The complexity of cooperation. Princeton, NJ: Princeton University Press.
- Basu, S., Dickhaut, J., Hecht, G., Towry, K., & Waymire, G. (2009). Recordkeeping alters economic history by promoting reciprocity. *Proceedings of the National Academy of Sciences of the U.S.A.*, 106(4), 1009–1014.
- Bednar, J., & Page, S. (2007). Can game(s) theory explain culture? The emergence of cultural behavior within multiple games. Rationality and Society, 19(1), 65–97.

- Bilancini, E., & Boncinelli, L. (2009). The co-evolution of cooperation and defection under local interaction and endogenous network formation. *Journal of Economic Behavior and Organization*, 70(1-2), 186–195.
- Bolton, G. E., Katok, E., & Ockenfels, A. (2005). Cooperation among strangers with limited information about reputation. *Journal of Public Economics*, 89(8), 1457–1468.
- 11. Bowles, S., & Gintis, H. (2011). A cooperative species: Human reciprocity and its evolution. Princeton, NJ: Princeton University Press.
- 12. Boyd, R., & Richerson, P. J. (1985). Culture and the evolutionary process. Chicago: University of Chicago Press.
- 13. Boyd, R., & Richerson, P. J. (2005). The origin and evolution of cultures. Oxford, UK: Oxford University Press.
- Brandt, H., Hauert, C., & Sigmund, K. (2003). Punishment and reputation in spatial public goods games. *Proceedings of the Royal Society B*, 270(1519), 1099–1104.
- Brandt, H., & Sigmund, K. (2005). Indirect reciprocity, image scoring, and moral hazard. Proceedings of the National Academy of Sciences of the U.S.A., 102(7), 2666–2670.
- Cong, R., Wu, B., Qiu, Y., & Wang, L. (2012). Evolution of cooperation driven by reputation-based migration. *PLoS ONE*, 7, e35776.
- 17. Dugatkin, L. A. (1997). The evolution of cooperation. BioScience, 47(6), 355-362.
- Dutton, W. H., Schneider, V., & Vedel, T. (2012). Ecologies of games shaping large technical systems: Cases from telecommunications to the Internet. In J. M. Bauer, A. Lang, & V. Schneider (Eds.), *Innovation policy and governance in high-tech industries: The complexity of coordination* (pp. 49–75). Berlin: Springer.
- Fletcher, J. A., & Doebeli, M. (2009). A simple and general explanation for the evolution of altruism. Proceedings of Royal Society B, 276(1654), 13–19.
- Fu, F., Wu, T., & Wang, L. (2008). Partner switching stabilizes cooperation in coevolutionary prisoner's dilemma. *Physical Review E*, 79(3), 036101.
- 21. Gigerenzer, G., Todd, P. M., & ABC Research Group. (1999). Simple heuristics that make us smart. Oxford, UK: Oxford University Press.
- Henrich, J. (2004). Cultural group selection, coevolutionary processes and large-scale cooperation. Journal of Economic Behavior and Organization, 53(1), 3–35.
- 23. Henrich, J. (2006). Cooperation, punishment, and the evolution of human institutions. *Science*, 312(5770), 60-61.
- Koenig, T., & Gogel, R. (1981). Interlocking corporate directorships as a social network. *American Journal of Economics and Sociology*, 40(1), 37–50.
- Kreps, D. M., & Wilson, R. (1982). Reputation and imperfect information. *Journal of Economic Theory*, 27(2), 253–279.
- Lehmann, L., & Keller, L. (2006). The evolution of cooperation and altruism: A general framework and a classification of models. *Journal of Evolutionary Biology*, 19(5), 1365–1376.
- Leimar, O., & Hammerstein, P. (2001). Evolution of cooperation through indirect reciprocity. Proceedings of the Royal Society B, 268(1468), 745–753.
- 28. Long, N. E. (1958). The local community as an ecology of games. *American Journal of Sociology*, 64, 251-261.
- Lubell, M. N., Robins, G., & Wang, P. (2011). Policy coordination in an ecology of water management games. Available at http://opensiuc.lib.siu.edu/pnconfs\_2011/22 (accessed March 2013).
- 30. Malerba, F. (2006). Innovation and the evolution of industries. *Journal of Evolutionary Economics*, 16(1-2), 3-23.
- Milinski, M., Semmann, D., Bakker, T. C. M., & Krambeck, H.-J. (2001). Cooperation through indirect reciprocity: Image scoring of standing strategy? *Proceedings of the Royal Society B*, 268(1484), 2495–2501.
- Milinski, M., Semmann, D., & Krambeck, H.-J. (2002). Reputation helps solve the "tragedy of the commons." *Nature*, 415(6870), 424–426.
- 33. Nelson, R. R., & Winter, S. G. (1982). An evolutionary theory of economic change. Cambridge, MA: Harvard University Press.

- 34. North, D. C. (1990). Institutions, institutional change and economic performance. Cambridge, UK: Cambridge University Press.
- Nowak, M. A., & Sigmund, K. (1998). Evolution of indirect reciprocity by image scoring. Nature, 393(6685), 573–577.
- 36. Nowak, M. A., & Sigmund, K. (2005). Evolution of indirect reciprocity. Nature, 437(7063), 1291-1298.
- Nowak, M. A., Tarnita, C. E., & Antal, T. (2010). Evolutionary dynamics in structured populations. *Philosophical Transactions of the Royal Society B*, 365(1537), 19–30.
- O'Hagan, S. B., & Green, M. B. (2004). Corporate knowledge transfer via interlocking directorates: A network analysis approach. *Geoforum*, 35(1), 127–139.
- Pacheco, J. M., Traulsen, A., & Nowak, M. A. (2006). Coevolution of strategy and structure in complex networks with dynamical linking. *Physical Review Letters*, 97(25), 258103.
- Panchanathan, K., & Boyd, R. (2003). A tale of two defectors: The importance of standing for evolution of indirect reciprocity. *Journal of Theoretical Biology*, 224(1), 115–126.
- 41. Pepper, J. W. (2007). Simple models of assortment through environmental feedback. Artificial Life, 13(1), 1–9.
- 42. Perc, M., & Szolnoki, A. (2010). Coevolutionary games: A mini review. Biosystems, 99(2), 109-125.
- Rand, D. G., Arbesman, S., & Christakis, N. A. (2011). Dynamic social networks promote cooperation in experiments with humans. *Proceedings of the National Academy of Sciences of the U.S.A.*, 108(48), 19193–19198.
- Santos, F. C., Pacheco, J. M., & Lenaerts, T. (2006). Cooperation prevails when individuals adjust their social ties. *PLoS Computational Biology*, 2(10), 1284–1291.
- Smaldino, P. E., & Lubell, M. (2011). An institutional mechanism for assortment in an ecology of games. PLoS ONE, 6(8), 23019.
- Smaldino, P. E., & Schank, J. C. (2012). Movement patterns, social dynamics, and the evolution of cooperation. *Theoretical Population Biology*, 82(1), 48–58.
- Suzuki, S., & Akiyama, E. (2005). Reputation and the evolution of cooperation in sizable groups. Proceedings of the Royal Society B, 272(1570), 1373–1377.
- 48. Szabó, G., & Fáth, G. (2007). Evolutionary games on graphs. Physics Reports, 446(4-6), 97-216.
- Tennie, C., Frith, U., & Frith, C. D. (2010). Reputation management in the age of the World-Wide Web. Trends in Cognitive Sciences, 14(11), 482–488.
- van den Bergh, J. C. J. M., & Gowdy, J. M. (2009). A group selection perspective on economic behavior, institutions and organizations. *Journal of Economic Behavior and Organization*, 72(1), 1–20.
- 51. Van Segbroeck, S., Santos, F. C., Pacheco, J. M., & Lenaerts, T. (2010). Coevolution of cooperation, response to adverse social ties and network structure. *Games*, 1(3), 317–337.
- 52. Wedekind, C., & Milinski, M. (2000). Cooperation through image scoring in humans. *Science*, 288(5467), 850–852.
- 53. Young, H. P. (1998). Individual strategy and social structure: An evolutionary theory of institutions. Princeton, NJ: Princeton University Press.