



Friendship, cliquishness, and the emergence of cooperation

Daniel J. Hruschka*, Joseph Henrich

Department of Anthropology, Emory University, 1557 Dickey Drive, Atlanta, GA 30322, USA

Received 16 December 2004; received in revised form 6 July 2005; accepted 7 July 2005

Abstract

The evolution of cooperation is a central problem in biology and the social sciences. While theoretical work using the iterated prisoner's dilemma (IPD) has shown that cooperation among non-kin can be sustained among reciprocal strategies (i.e. tit-for-tat), these results are sensitive to errors in strategy execution, cyclical invasions by free riders, and the specific ecology of strategies. Moreover, the IPD assumes that a strategy's probability of playing the PD game with other individuals is independent of the decisions made by others. Here, we remove the assumption of independent pairing by studying a more plausible cooperative dilemma in which players can preferentially interact with a *limited* set of known partners and also deploy longer-term accounting strategies that can counteract the effects of random errors. We show that cooperative strategies readily emerge *and* persist in a range of noisy environments, with successful cooperative strategies (henceforth, *cliquers*) maintaining medium-term memories for partners and low thresholds for acceptable cooperation (i.e. forgiveness). The success of these strategies relies on their *cliquishness*—a propensity to defect with strangers if they already have an adequate number of partners. Notably, this combination of medium-term accounting, forgiveness, and cliquishness fits with empirical studies of friendship and other long-term relationships among humans. © 2005 Published by Elsevier Ltd.

Keywords: Iterated prisoner's dilemma; Partner choice; Friendship; Altruism; Cooperation

1. Introduction

The emergence of cooperation among unrelated individuals poses a continuing challenge to theories in biology and economics (Hammerstein, 2003). Much research on cooperation has focused on how natural selection might favor strategies that cooperate in situations analogous to the two-person prisoner's dilemma (PD). In the two-person PD, two individuals have the opportunity to exchange some favor but face the possibility of exploitation by their partner. When the individual benefit of cooperating (e.g. $B = 200$) is greater than the cost (assume $C = 100$), then the overall payoff when both cooperate ($2 \times (B - C) = 200$) is greater

than if neither does (zero). An evolutionary challenge arises because there is always a possibility that one partner will not cooperate, thereby taking the benefit (B) and leaving the partner with less ($-C$) than if he or she had not given anything (Again, zero).¹ In the purely one-shot game among non-relatives, the only evolutionarily stable solution is to defect.

Trivers (1971) suggested that individuals might achieve cooperation if they have the opportunity to exchange cooperation over a sequence of moves. This idea was later formalized and modeled by Axelrod and Hamilton (Axelrod, 1984; Axelrod and Hamilton, 1981). From this early work, the iterated prisoner's dilemma (IPD) emerged as the canonical model for studying the

Abbreviations: B/C, Benefit-to-cost of cooperating; IPD, Iterated prisoner's dilemma; Tft, Tit-for-tat; Oft, Out-for-tat

*Corresponding author. Tel.: +1404 636 0144; fax: +1404 727 2860.

E-mail addresses: dhrusch@sph.emory.edu (D.J. Hruschka), jhenrich@emory.edu (J. Henrich).

¹Based on the two partner's choices, there are four possible payoffs: $R = B - C$ if both cooperate, $T = B$ if your partner cooperates but you defect, $S = -C$ if you cooperate and your partner defects, and $P = 0$ if neither cooperates. Given that B/C is greater than 1, we know that $T > R > P > S$, and $2R > T + S$.

evolution of cooperation via reciprocity in repeated interactions. To model long-term relationships, the evolutionary IPD randomly pairs individuals and then forces them to play a repeated PD, where the probability of play continuing to the next interaction (after the first) is fixed.² In such a setup, reciprocally cooperative strategies can resist invasion by pure defectors—if initially common—and can often maintain high levels of cooperation. This early work suggests that successful reciprocating strategies should be NICE (always cooperate initially) and PROVOKABLE (withdraw cooperation at the first sign of defection), qualities associated with the oft-cited strategy of tit-for-tat (TFT) (Bendor and Swistak, 1997; Gotts et al., 2003).

Later work, however, revealed both theoretical and empirical challenges to the early success of reciprocity-based strategies—especially TFT—in the IPD. First, various forms of errors or noise can devastate such reciprocal cooperation. Both *errors of implementation*, which occur when individuals mistakenly apply their strategies (e.g. defect when they mean to cooperate), and *errors of perception*, which occur when individuals misperceive their partner's actions (e.g. believe their partner defected when he actually cooperated), reduce the ability of reciprocal strategies to maintain cooperation (Bendor et al., 1996; Boerlijst et al., 1997). Second, the success of particular strategies depends critically on the distribution or 'ecology' of other strategies in the population which may change through mutation and selection (Boyd and Lorberbaum, 1987; Lorberbaum, 1994; Lorberbaum et al., 2002). And finally, there remains the long-recognized problem that reciprocal strategies generally cannot spread when rare without some form of ordered spatial or interactional distribution (Axelrod, 1984).

A number of empirical concerns accompany these theoretical shortcomings, particularly in the human case. First, by contrast with the canonical IPD, individuals in many real situations are often neither locked into a sequence of interactions with unwanted partners nor forced to switch randomly to a new partner after some number of iterations. Numerous examples from humans, and other animals, show that individuals often leave relationships with one partner to find others with whom they might have better payoffs (Bshary and

Noe, 2003; Hammerstein, 2003; Kelley and Thibaut, 1978; Noe and Hammerstein, 1994). Analytical IPD models typically assume that strategies meet with a probability equal to their current representation in the population. This would not be the case, however, if partners who achieve successful cooperation tend to stick together longer than partners who do not. Second, while interactions among strangers or acquaintances seem empirically consistent with TFT-like strategies (maintain close accounting and remain highly provokable), interactions within long-term human relationships such as friendships do not. Empirical evidence indicates that trusted partners do not closely track particular exchanges and are not easily provoked into defection by a partner's occasional defection (Silk, 2003). Moreover, humans are more willing to engage in costly cooperation with trusted partners than with strangers, and thus are not generally NICE in the way that successful strategies often are in the IPD (Markovits et al., 2003).

An alternative formulation for the evolution of pairwise cooperation—the explicit partner choice model—addresses several of the theoretical challenges just described for the IPD while yielding results that are substantially more consistent with the existing empirical work on humans than are those from the IPD. In explicit partner choice models, players are able to influence their pairings (and thereby the pairing of others) using past experience. Thus, intuitively, they do not get 'stuck' in a sequence of interactions with particular individuals (or strategies). Such models provide added opportunities for avoiding poor-quality partners, while forcing individuals to weigh new trade-offs between search costs and the opportunity costs of staying with particular partners (Connor, 1992; Enquist and Leimar, 1993; Friedman and Hammerstein, 1991; Noe and Hammerstein, 1994). Existing work on partner-choice models suggests that they have a number of advantages over the canonical formulation (Batali and Kitcher, 1995; de Vos et al., 2001; de Vos and Zeggelink, 1997; Hauk, 2001; Hayashi and Yamagishi, 1998; Peck, 1993; Sherratt and Roberts, 1998; Skyrms and Pemantle, 2000; Stanley et al., 1994). First, when partner choice is permitted, strategies can often rapidly achieve high levels of cooperation, even when the cooperative strategies are not initially common (Ashlock et al., 1996). Second, partner choice may resolve the problem of cyclical invasions in a changing ecology of strategies (de Vos et al., 2001; Sherratt and Roberts, 1998). Third, the basic setup that allows individuals to select and/or reject their partners has an intuitive plausibility for humans that the canonical 'lock-in' models of reciprocal altruism lack. Consistent with this approach, non-evolutionary tournament simulations suggest that partner choice algorithms, such as out-for-tat (OFT), may be more important in facilitating

²Another interpretation of the iterated IPD is that individuals play the PD with every other individual in the population for a specified number of rounds. In this sense, the IPD allows individuals a kind of partner choice by allowing them to cooperate with some individuals but not with others (if the payoff to mutual defection is equivalent to the payoff of non-interaction). This, however, forces each player to expend time in an interaction with each member of the population, rather than permitting players to focus their efforts on a few known, mutually cooperative partners. Conversely, in our model, strategies that develop cooperative relationships effectively remove themselves from the mix of strategies that defectors, for example, might encounter and exploit.

cooperation than are decisions to cooperate once partners have been paired (Yamagishi et al., 1994).

Despite the promise of the explicit partner choice approach, the existing work has some shortcomings. First, previous research has ignored the potential effects of errors, which have been shown to be critically important in the canonical model (Bendor et al., 1996). Second, and more importantly, previous approaches assume that players maintain complexly integrated memories of past interactions with *every other* individual in a group, making the models implausible for human populations of more than 5–10 individuals (Dunbar, 1998; Milardo, 1992). Finally, the existing simulations have been uniformly limited to fewer than 60 players (often using only 20), restricting the generalizability of findings to more realistic populations and raising concerns about the effects of random drift (Leimar and Hammerstein, 2001).

In integrating and extending existing work on partner choice models, our simulation model focuses on five key aspects. First, we synthesize prior strategies by capturing a substantial portion of the strategy-space with 4 simultaneously evolving parameters (1 dichotomous and 3 continuous) and by permitting mutation to change these parameters. Second, we include two types of error, one in which individuals mistakenly defect when they mean to cooperate (and vice versa) and the other in which they are unable to meet with a preferred partner. Third, we restrict individuals' memories to one number for K different partners. For much of the analysis below, K is set to three, as this is cognitively plausible for many primates and is in a range observed for the number of close, supportive ties observed in humans (Dunbar, 1998; Milardo, 1992; Sugawara, 1984). We also briefly examine the effect of changing K , varying it from one to eight. Third, we consider the dynamics of strategies in much larger populations than those considered in past simulations ($N = 100$ –2000), which cover a more representative set of common group sizes found among extant human populations (Dunbar, 1998). Fourth, our model is fully evolutionary. We have avoided non-evolutionary pairwise tournaments and included mutation on all of our strategy parameters in order to avoid the restricted ecologies that have clouded prior work. Finally, the model permits us to examine prior assumptions about relationship decay in the IPD environment and to show the assumption of geometric decay is only valid in the long-run and not during initial interactions (Feldman and Thomas, 1987; Getty, 1987).

2. The model and simulation

To explore the evolution of cooperation, we constructed a computer simulation in which individuals are paired to play one round of the PD game according to

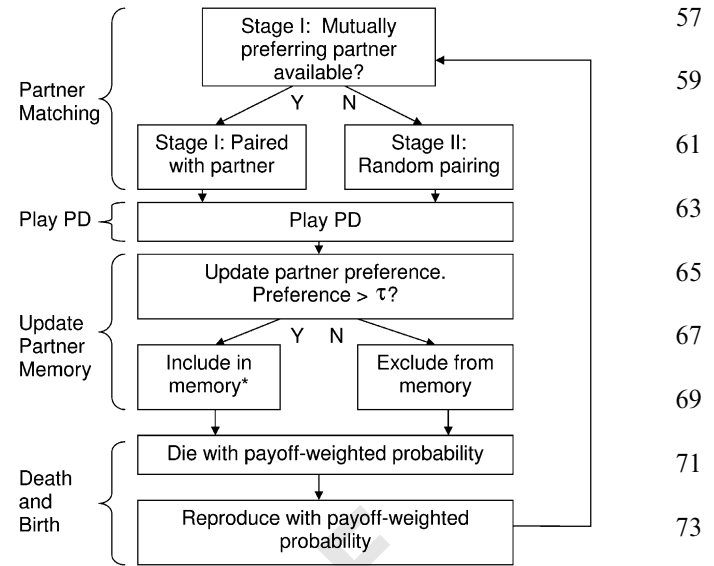


Fig. 1. Simulation flowchart. The asterisk in the “Include in Memory” box indicates that a new partner is added to memory contingent on (1) the available memory slots and (2) the ranking of its preference value relative to existing partners.

mutual preferences for each other.³ After each game, players update their preferences for specific partners based on their payoff from the current round. Next, individuals die and reproduce based on their payoffs in the game, and then individuals are paired again for a new round of the PD game based on their updated partner preferences (see Fig. 1 for an outline of the simulation process).

We consider a population of fixed size ($N = 100, 400, 800$, or 2000) in which each individual maintains a strategy that specifies when to cooperate in the PD and how to calculate partner preferences.⁴ In each round, individuals apply this strategy to a PD game with a fixed payoff matrix. We explored PD payoff matrices with benefit-to-cost ratios (B/C) varying from 1 to 16. Table 1 represents two examples of these payoff matrices, one for a high ratio ($B/C = 8$) and one with a lower ratio ($B/C = 2$) when $C = 100$.

2.1. Memory, partner preferences, and matching

For a limited number of potential partners, individuals can remember a single real number in the interval $[0,1)$ that summarizes past interactions with each of

³The simulation was programmed in Visual Basic 6.0 and is available from DJH (dhrusch@sph.emory.edu).

⁴By using one population (of size N), we have potentially conflated the effects of (a) drift with (b) the number of potential partners to which one has access. However, our results largely alleviate this concern, as the effects of varying N look remarkably like the anticipated effects of drift, with the influence of N disappearing as it gets large.

Table 1
Payoffs to partner A

	$B/C = 2$		$B/C = 8$	
	B cooperates	B defects	B cooperates	B defects
A cooperates	200	0	800	0
A defects	300	100	900	100

these partners and represents the strength of preferences for future matching with each of the partners. As noted, this number is updated after each play of the PD game. Each strategy determines how this index is recalculated and when to add and remove old partners from its list.

Matching occurs in two stages: (I) the *preference stage* in which individuals attempt to pair with preferred partners and (II) the *random assignment stage* in which those who have not successfully matched with a preferred partner are randomly assigned among themselves. In the preference stage, an individual (referred to as ego) is chosen randomly from the population for matching. Subsequently, the matching algorithm:

- (1) Considers ego's most preferred partner (of the K partners that are potentially stored in ego's memory). If that partner also prefers ego, and that partner is not yet paired, then ego and the partner are paired (with some probability of failure, e_m). If no partner exists in ego's memory, then ego is moved to stage II (placed in the pool for random pairing).
- (2) If the first preferred partner is not available then the algorithm considers the next preference, and so on.
- (3) If none of the K preferred partners are available (i.e. the partners have already been matched, they do not prefer ego, or they are unable to be paired due to matching error), then ego moves to stage II.
- (4) Another unpaired individual is randomly chosen and the matching routine repeats, starting at step 1.

After all individuals' preferences have been considered, stage II randomly pairs the remaining individuals who were unable to match with preferred partners.

2.2. Decisions to cooperate

Once all individuals are paired, they decide to cooperate or not cooperate with their partner based on two parameters (D and χ). An unconditional defector ($D = 1$) always defects (with probability $1 - e_s$). A cooperative strategy ($D = 0$) cooperates if it has been matched with a preferred partner (with probability $1 - e_s$, e_s is the 'strategy error' and is detailed below). If it is not able to pair with any of its preferred strategies, then its choice to cooperate is contingent on whether it already has K preferred partners. If its memory is full, then it

defects with its randomly assigned partner with probability χ —thus the evolving parameter χ measures NICENESS. However, if the strategy still has memory available, then it cooperates.⁵

2.3. Updating partner preferences

Once all strategies have made their decision to cooperate or defect, everyone updates their partner-memory based on two parameters: a memory weight placed on past interactions (δ) and a threshold for acceptable cooperation (τ). Each strategy calculates future preferences for its current partner by re-adjusting past preferences with payoffs from the most recent round:

$$pref_{ijk} = \delta pref_{ij(k-1)} + (1 - \delta) M_{ijk}.$$

Here, $pref_{ijk}$ is the preference that player i has for player j on the k th round. M_{ijk} is the payoff to player i from the k th round of play with player j (if they were paired) divided by the maximum possible payoff in the PD game. For newly matched individuals, $pref_{ij0} = 0$. If no pairing occurs in a particular round between i and j , $pref_{ijk}$ is not updated. This accounting method has several useful properties:

1. It is simple and requires keeping only one number in long-term memory.
2. It integrates past and present experience.
3. If δ is non-zero, it weights the long-term expected value of M_{ijk} by the number of times the two players have interacted together (m). Specifically, under conditions where the payoffs from a partner are temporally homogenous (as they are for most strategies considered here), the expected value of $pref_{ijm}$ is equal to $E(M_{ij.})(1 - \delta^m)$ —the product of the expected value of past payoffs $E(M_{ij.})$ and a monotonically increasing function of the duration of the relationship $(1 - \delta^m)$.⁶

This accounting scheme is similar to the weighted average of past interactions proposed by Bendor (1987). Unlike Bendor's approach, however, the strategies in this paper do not base their decisions to cooperate on

⁵Below, we also explore the situation in which strategies could use χ to regulate defection in both the random matching pool and with mutually preferred partners. We present results of these simulations in the section titled, 'Part-Time Defectors versus Cliques'.

⁶Derivation of $E(pref_{ijm}) = E(M_{ij.})(1 - \delta^m)$. Let S be the entire set of possible temporal sequences for the payoffs (M_{ijk}) between individual i and j . Also, let $E(M_{ijk})$ be the expected value of M_{ijk} over S . If M_{ijk} is temporally homogenous, then the expected payoff at one interaction will be the same as the expected payoff at another interaction (i.e. $E(M_{ijk}) = E(M_{ijk}) = E(M_{ij.})$). So $E(\sum_{k=1}^m M_{ijk} \delta^{m-k} (1 - \delta)) = \sum_{k=1}^m E(M_{ij.}) \delta^{m-k} (1 - \delta) = E(M_{ij.}) \sum_{k=1}^m \delta^{m-k} = E(M_{ij.})(1 - \delta^m)$.

this index. Rather, they use the index to choose partners for play.⁷

If ego and a new partner cooperate, ego may decide to include the partner in his memory. If ego has a free memory slot, then he adds the new partner to his memory. However, if ego's memory is full, then he compares his preference for the new partner with his preference for existing partners. If his preference for the new player is greater than that for at least one of ego's existing partners, then ego replaces the partner having the lowest preference score with the new partner's information.

Each strategy uses a threshold for acceptable cooperation (τ) to remove uncooperative partners from memory. Partners whose value $pref_{ijk}$ drops below this threshold are removed from ego's memory, while those with $pref_{ijk} > \tau$ are ranked from highest to lowest—keep in mind that players maintain a maximum of K preference values in mind at any one time.

2.4. Life, death and reproduction

Unlike most implementations of the evolutionary IPD, fitness-biased reproduction occurs after each play of a one-shot game. Fitness for an individual is based on payoffs from the most recent game *plus* some carry-over from past games:

$$F_{jk} = \pi_{jk} + \alpha F_{j(k-1)}.$$

Here, F_{jk} is the fitness of individual j after round k . The payoff, π_{jk} , is the payoff that individual j received from play in his k th game. The quantity $\alpha F_{j(k-1)}$ is the fitness carried over from the previous round, so α is the fraction of fitness not used up between rounds. In this paper, we present results for $\alpha = 0.6$, although basic results do not differ for $\alpha = 0$ (no carryover) or 0.30.

This form of fitness accounting captures two key aspects of fitness. First, it creates a geometric decay that prevents individuals from accumulating unrealistically large fitness levels relative to others. This is consistent with the observation that there are biological and social upper limits on one's ability to 'store' bases for fitness (e.g. food storage, rank achievement, and energy stores). Our form sets a maximum limit, $\max(F_j)$, on fitness:

$$\max(F_{jk}) = \frac{B}{1-\alpha}.$$

Here, B is the benefit to cooperation. Second, as one approaches this ceiling, there is generally decreasing marginal capacity for the bases of fitness (e.g. energy stores) to produce survival and reproduction. This fitness formulation guarantees a monotonically decreasing slope in the curves linking (1) accumulated fitness

and survival and (2) accumulated fitness and reproduction.

Selection occurs after each round with the elimination of one individual based on a fitness-weighted probability. An individual's probability of elimination (Pd_{jk}) is

$$Pd_{jk} = \frac{G_{jk}}{\sum_{All\ j} G_{jk}},$$

where

$$G_{jk} = 1 - \frac{F_{jk} - \min(F_{.k})}{\max(F_{.k}) - \min(F_{.k})}.$$

Upon selection, knowledge of this individual is also erased from the memories of other individuals.

Similarly, on the reproduction side, one of the remaining individuals is selected to reproduce and fill the empty slot based on a fitness-weighted probability (Pr_{jk}):

$$Pr_{jk} = \frac{F_{jk}}{\sum_{All\ j} F_{jk}}.$$

A *generation* is defined as the number of rounds it takes to reproduce the number of organisms found in the population. For example, a generation in a population of 100 players is 100 rounds.⁸

2.5. Parameterizing the strategies

As described above, four variables parameterize all possible strategies that we study:

- (1) *Always defect*, D : A dichotomous parameter indicating whether a strategy unconditionally defects ($D = 1$) or follows an algorithm determined by the next three variables ($D = 0$).
- (2) *Past weighting*, δ ($0 \leq \delta \leq 1$): The relative weight of past experience in determining partner preferences. If δ is near one, current preference for a partner is heavily weighted toward one's prior preference. If $\delta = 0$ then only the payoff from the most recent round matters, as is the case for TFT or OFT.
- (3) *Threshold of acceptable cooperation*, τ ($0 \leq \tau \leq 1$): The threshold used to judge the acceptability of a partner's past cooperation. If a player's preference for a partner drops below τ , the partner is removed from the player's memory.
- (4) *Cliquishness* χ ($0 \leq \chi \leq 1$): The probability that a player using this strategy will defect with a new partner given that she already maintains K acceptable partners.

⁸While non-overlapping generations would allow a more direct comparison with the canonical IPD, our approach, because we are interested in the emergence of partner networks that outlast any given individual, demands overlapping generations.

⁷Whether unconditional defectors have memory for partners or not is a moot point because by defecting they are unable to make themselves preferred by another partner.

Table 2
Common strategies mapped onto parameter specifications

Strategy	Past weighting(δ)	Provokability(δ)	Cliquishness(χ)	D
Pure defector	n/a	n/a	n/a	1
Pure cooperator	Any	Any	0	0
Pure cooperator (no memory)	Any	$\tau > 1 - \delta$	0	0
Out-for-tat	0	High	0	0
Cliquer	Moderate	Low	High	0

Different specifications of the four variables (δ , τ , χ , D) correspond to the commonly known strategies from the canonical IPD—see Table 2. While unconditional defectors are exactly the same as the analogous strategies used in canonical models, pure cooperator and analogs to TFT strategies require some comment. Pure cooperators here always intend to cooperate, but they can still exercise partner preferences, and thus some cooperators are more discriminating than are others. When $\tau > 1 - \delta$, however, cooperative strategies are unable to register partner preferences and so act like pure cooperators without partner memory. The classical definition of TFT strategies relies on the forced nature of interactions in the canonical IPD, and there are several ways to translate the NICE and PROVOKABLE properties of TFT into partner choice scenarios. In our simulation, possible parameter combinations include out-for-tat (OFT) strategies that *leave* defecting partners—an approach achieving higher payoffs in tournament studies than traditional TFT strategies that stay with a partner and defect on the next round (Yamagishi et al., 1994). Like pure cooperators, OFTERS have partner preferences, but those preferences are based entirely on what happened the last time those two strategies met. In addition, OFTERS are always NICE (initially cooperative) to individuals not contained in their K memory slots. Table 2 includes a novel strategy, *cliquer*, which we derived based on empirical studies of human friendships.

Cliquer strategies have a number of characteristics worth noting, some of which contrast with TFT.

- (1) A high δ means that cliquers place value on the duration of the relationship. They prefer partners with whom they have developed a long series of beneficial interactions.
- (2) A high δ combined with a low τ makes them insensitive to short-term aberrations, thus making both cooperation and an individual's preferences robust against any single defection (or even a few defections) after a long, cooperative relationship. cliquers, however, remain sensitive to a series of defections—as do humans (Silk, 2003).
- (3) *Conversely*, these same preferences make individuals more sensitive to a partner's behavior *early* in the relationship—as in humans.
- (4) Unlike TFT, *cliquers* are not generally NICE. Cliquers

cooperate with preferred partners, but cooperate with novel partners only if their memory slots are not full. Faced with a new partner when it has enough acceptable partners, a cliquer defects.

We will refer to the broad class of strategies where $D = 0$ (cliquer, cooperator, and OFT) as 'cooperative strategies' (CS) in contrast to the specific sub-class of 'pure cooperators' ($\chi = 0$ and $D = 0$).

2.6. Mutation

When a strategy is selected for reproduction, its four parameters are reproduced with mutational change. The offspring's values for δ , τ , χ are mutated from the parent's value based on a normal distribution with mean zero and standard deviation, σ (typically set to 0.05), and the value for D is reversed with a probability of p_{switch} (typically set to 0.05).⁹ In cases where mutations produce a value outside defined conditions (e.g. $\tau < 0$) the value is set to the nearest boundary value (e.g. $\tau = 0$). When a defector mutates into a cooperative strategy, the three strategy parameters, δ , τ , and χ are randomly assigned (uniform distribution) to values in the interval [0,1).

To address concerns about this random reassignment, we also considered an alternative approach to mutation in which all strategies (even unconditional defectors) maintained strategy parameters, δ , τ , and χ but only those strategies with $D = 0$ used their parameters. In this way, when an unconditional defector mutated into a cooperative strategy, new values for δ , τ , and χ did not need to be randomly assigned. Simulation results using the two approaches to mutation did not qualitatively differ.

2.7. Errors of execution and matching

We examine the effects of two kinds of errors. Individuals in our model execute their intended behavior

⁹We also examined simulations at lower values of σ and p_{switch} (0.01 and 0.005), and found that cooperative strategies achieved even higher proportions in the population (though after a much greater number of generations) when they held a selective advantage over unconditional defectors. For practical considerations, we selected these higher mutation rates, but they do not substantially change our findings.

(based on their strategy) with probability $1-e_s$, and execute the opposite behavior with probability e_s . Thus, if an individual plans to cooperate with a preferred partner, she will ‘accidentally’ defect a fraction e_s of the time. We explore these *strategy errors* from $e_s = 0.00$ to 0.20 . We also consider errors in partner choice or *matching errors*. Specifically, if one has identified a mutually preferring partner in the matching round, there is a probability e_m (ranging from 0.00 to 0.70 depending on the simulation) that one will not successfully pair with that player. If this occurs, then one must attempt to pair with the next preferred partner. In these simulations, we do not consider errors of perception (Boerlijst et al., 1997).

3. Results

In using simulations to examine the emergence and persistence of cooperation, we focus on (1) the relative selective advantage of strategies in a range of environments (simulation sets 1 and 3), and (2) the ability of an initially rare strategy (e.g. cliquer or OFT) to invade a population of alternative strategies (e.g. unconditional defectors) (simulation sets 2 and 4).

3.1. Selective advantage in a range of environments (simulation set 1)

To assess selective advantage under mutation, we first examine the relative frequencies of different strategies under a range of environmental conditions (i.e. by varying B/C , e_s , e_m , and N). We define the *selective advantage* of cooperative strategies ($D = 0$) as the extent to which their population prevalence deviates from that expected under neutral selection (i.e. 50%). Specifically, at reproduction, D has the same probability of switching from 0 to 1 as it does from 1 to 0. Thus, if the two variants of D (0 and 1) experience equal selective pressures, then we would expect values of this parameter to reach equal proportions in the population. However, if there are consistently more strategies with $D = 0$, then we conclude that cooperative strategies are favored by selection. In most simulations, the proportion of cooperative strategies (denoted as \overline{CS}) quickly settles into a confined area of the parameter space. Consequently, we used the mean proportion of cooperative strategies from between 100 and 200 generations ($200 \text{ generations} = 200 \times N \text{ rounds}$) to assess the selective advantage of cooperative versus unconditional defector strategies.

The same rationale was applied to determine which strategy parameter values (δ , τ , χ) were favored by selection. Specifically, under neutral selection with mutation, the expected population mean for a parameter is 0.50 and deviations from this value indicate that

selection is acting on the parameter values. In addition, a tighter parameter variance then observed in the neutral model indicates selection is operating on the parameter.

For $N = 100$ and 400 , we examined the sensitivity of strategy parameters (D , δ , τ , χ) to changes in environmental conditions ($B/C \sim 1-16$, $e_s \sim 0.00-0.20$, $e_m \sim 0.00-0.70$). For each scenario, we ran 5 trials of 200 generations to examine the consistency of results. To explore the effect of population size (and drift) on the simulations we conducted single runs at $N = 800$ and 2000 for 100 generations in four environmental conditions ($B/C = 4, 16$; $e_s = 0.05, 0.15$, $e_m = 0.05$).

3.1.1. Findings: selective advantage of cooperative strategies

Across a wide range of environments, CS —with different values of the evolving parameters—emerged and sustained mean proportions in the population well above 50% (the expected proportion under neutral selection) with much smaller temporal variation than observed in the null model. Fig. 2A illustrates the temporal dynamics of a typical run, while Fig. 2B shows that same run without selection. The curve marked CS gives the proportion of cooperative strategies in the current population across 200 generations ($200 \times 400 = 80,000$ rounds). As expected, parameter means in the model without selection all exhibit a great deal of temporal variation with means around 0.5.

The prevalence of cooperative strategies (\overline{CS}) is sensitive to the benefit-to-cost ratio (B/C), the population size (N), the strategy error (e_s), and the matching error (e_m)—see Fig. 3. First, there is a precipitous decline in the frequency of cooperative strategies between $B/C = 4$ and 1, which is consistent with the benefit-to-cost thresholds derived from analytical models. Second, increasing the strategy error (e_s) decreases the mean proportion of cooperative strategies in the population (Fig. 3B). This effect is most pronounced in smaller populations with lower benefit-to-cost ratios. In these situations (e.g. $N = 100$, $B/C = 2$), cooperative strategies lose their selective advantage at high rates of strategy error ($e_s = 0.20$). By contrast, cooperative strategies in large populations with high B/C ratios are able to withstand quite high rates of e_s .¹⁰ As with strategy errors, Fig. 3C shows that increasing errors in matching with a preferred partner (e_m) decreases the mean proportion of cooperative strategies in a popula-

¹⁰Figs. 3A and B both indicate that larger population sizes facilitate the success of cooperative strategies (CS) against defectors. The prevalence of CS in larger populations ($N = 800$ and 2000) is nearly identical to that in $N = 400$, indicating that the effect of population size is most pronounced when moving from $N = 100$ to 400 and then decreases rapidly toward zero beyond $N = 400$. This pattern of effects suggests increasing N reduces the influence of random drift, which acts against CS in of many these simulations.

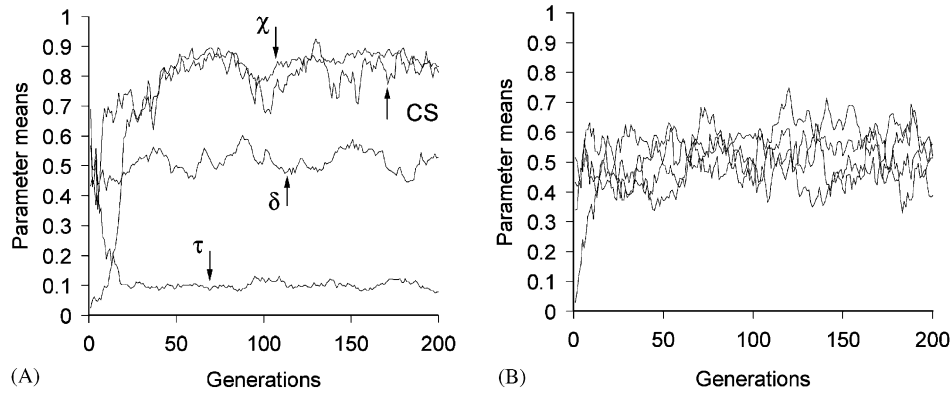


Fig. 2. Evolution of strategy parameters ($N = 400$, e_s and $e_m = 0.05$, $B/C = 2$). (A) Within-run parameter means under selection: $\overline{CS} = 0.81$ (SD = 0.05), $\bar{\delta} = 0.51$ (SD = 0.03), $\bar{\tau} = 0.10$ (SD = 0.01), $\bar{\chi} = 0.85$ (SD = 0.02). (B) Under neutral selection, within-run parameter means approximate 0.5 (0.48–0.54) with larger standard deviation (SD = 0.06–0.09). SD is an index of temporal variation in population means.




Fig. 3. Prevalence of cooperative strategies. (A) By B/C ratio (e_s and $e_m = 0.05$). (B) By strategy error ($e_m = 0.05$). (C) By partner matching error ($e_s = 0.05$). Data points represent 5 simulations of 200 generations, and means are taken from generations 100–200. Error bars represent standard deviations of the mean of population means taken from 5 runs. Values for $N = 800$ and 2000 were too close to $N = 400$ to represent on the graph. For example, in (A), the values for $N = 800$ at $B/C = 2, 4$ and 8 were $0.84, 0.95$ and 0.97 , respectively.

tion, and the most pronounced effect is observed in small populations with lower B/C ratios. The large standard deviation for the mean proportion of cooperative strategies where $N = 100$, $B/C = 2$, $e_s = 0.05$, and $e_m = 0.30$, results from the long-term oscillation of the population between two regimes—one dominated by unconditional defectors and one with a majority of cooperative strategies.

3.1.2. The selective advantage of strategy parameters

After an initial period of change, the continuous parameters for cooperative strategies (δ , τ , χ) were each distributed, usually rather tightly, around one central tendency within a population. In addition, the parameter variances were much smaller than those observed in the neutral model. This suggests that the three continuous strategy parameters were under considerable selective pressure, and indicates that the parameters of

CS can be accurately summarized using their population means.

The specific parameters of successful cooperative strategies depends on the environmental details:

Cliquishness (χ): In those environments where cooperative strategies maintain a selective advantage, mean cliquishness consistently has values greater than 0.6 with most values ranging between 0.8 and 0.9. With increasing B/C ratios, the long-term mean¹¹ values for cliquishness (χ) decrease, suggesting that cliquishness among cooperative strategies is most favored by selection when B/C ratios are low (Fig. 4A). Neither errors (e_s or e_m) nor population size (N) have substantial impacts on cliquishness, although in extremely large

¹¹Long-term means taken from generations 100–200 under the following condition: $B/C = 1, 2, 3, 4, 8, 16$; strategy or matching error rates = 0.00, 0.01, 0.05, 0.10, 0.15, 0.20; $N = 100, 400$.




Fig. 4. Mean strategy parameters (χ , δ , and τ). (A) By B/C ratio ($N = 400$, e_s and $e_m = 0.05$). (B) By strategy error ($N = 400$, $B/C = 4$, $e_m = 0.05$). (C) By matching error ($N = 400$, $B/C = 4$, $e_s = 0.05$). (D) By population size ($B/C = 4$, e_s and $e_m = 0.05$). Data points represent 5 simulations of 200 generations, and means are taken from generations 100–200. Error bars represent standard deviations of the mean of population means taken from 5 runs.

populations ($N = 2000$) high cliquishness appears to experience weaker selection.

Past weighting (δ): The mean weight of past experience, δ , increases as both B/C and e_s increase (Figs. 4A and B). The increase in mean δ is slight for B/C , rising from 0.51 at $B/C = 2$ to 0.60 at $B/C = 16$. For strategy error, there is a dramatic increase in mean δ from $e_s = 0$ to $e_s = 0.01$, with a slower rise from $e_s = 0.01$ to $e_s = 0.20$. This substantial increase suggests that strategies that place more weight on past interactions succeed most in environments with higher rates of execution errors. Conversely, in environments characterized by no execution error, strategies with low δ (i.e. similar to OFT to TFT) have a selective advantage. Population size also influences the long-term mean of δ , with increasing N leading to greater weighting of past experience.

Because δ values are often in the vicinity of 0.50 (i.e. values also favored by neutral selection) we compared the standard deviations seen in Fig. 4A and B with those from our neutral model simulations and found that these parameters are clearly under strong selection. Not only are most δ values significantly different from 0.5, but the within-run standard deviation for δ under neutral selection ($SD = 0.07$) are two to three times greater than that for δ under selection in any environmental condition ($SD = 0.02$ – 0.03).

Threshold of acceptable cooperation (τ): In those environments where CS maintain a selective advantage, mean τ are consistently lower than 0.2. Increasing values of strategy error lowers the long-term mean values of τ , with the most dramatic decline occurring between $e_s = 0$ and 0.01. This suggests that at larger values of strategy error, more forgiving strategies succeed. Neither matching error (e_m) nor benefit-to-cost ratio has an impact on long-term means of τ . With increasing population size, long-term means of τ also decrease slightly.

The primary advantage of high τ values is the ability to leave partners who start by cooperating but then begin to defect consistently after some period of

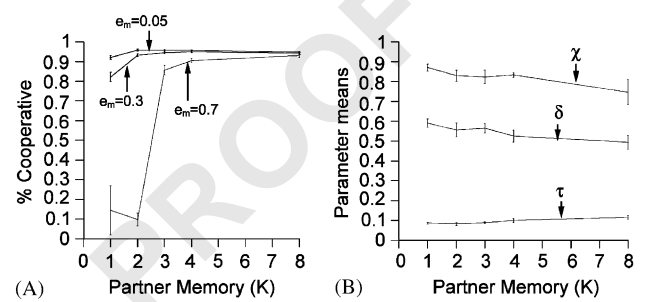


Fig. 5. The effects of partner memory and matching error. (A) on the prevalence of cooperative strategies ($N = 400$, $B/C = 4$, $e_s = 0.05$). (B) on mean CS parameters ($N = 400$, $B/C = 4$, e_s and $e_m = 0.05$). Data points represent 5 simulations of 200 generations, and means are taken from generations 100–200. Error bars represent standard deviations of the mean of population means taken from 5 runs.

cooperation. On the other hand, there is a strong selective force pushing τ towards 0—the need to forgive long-term cooperative partners who make a consecutive string of accidental defections. In situations where defectors are more sophisticated and can change their rates of defection, we would expect τ to increase.

3.2. Selective advantage when varying partner memory

The above analyses assume strategies can have a maximum of $K = 3$ preferred partners in memory. To assess the sensitivity of results to this assumption, we examined the same population dynamics as in the previous section while varying K ($= 1, 2, 3, 4, 8$). We also varied matching error ($e_m = 0.05, 0.30, 0.70$), expecting that with increasing matching error, larger K would be necessary to buffer against the possibility of missed pairings. Fig. 5A shows that with low matching error ($e_m = 0.05$), CS with memory slots for 2 partners have only a slightly greater selective advantage over pure defectors than do CS with only one memory slot. However, matching error increases this effect, and in

Table 3
Probability of cliquer and OFT invading and displacing a population of unconditional defectors (out of 10 trials, $N = 400$, $e_m = 0.05$)

Starting proportion of CS								
Invading strategy	e_s	B/C	0.005	0.01	0.02	0.05	0.10	0.30
Cliquer	0.05	2	0.2	0.6	0.9	1	1	1
		4	0.6	1	1	1	1	1
	0.00	2	1	1	1	1	1	1
		4	1	1	1	1	1	1
OFT	0.05	2	0	0	0	0	0	0
		4	0	0	0	0	0	0
	0.00	2	0.45	0.40	0.38	0.40	0.43	0.43
			(0.07) ^a	(0.06)	(0.08)	(0.08)	(0.05)	(0.09)
		4	0.55	0.56	0.58	0.55	0.55	0.57
			(0.02)	(0.03)	(0.03)	(0.04)	(0.04)	(0.03)

^aIn cases where OFT did not displace unconditional defectors, but reached stable proportions in a population, the 10-run mean and (standard deviation) of final proportions for OFT is listed.

environments with high matching error ($e_m = 0.70$), CS with memory slots for fewer than three partners do not have a selective advantage over unconditional defectors. We would expect this interaction between K and e_m because the probability of failing to match with *all* of K partners (e_m^K) depends on both the number of partners (K) and the probability of failing to match with each partner independently (e_m).

Fig. 5B shows that the effect of K on CS parameters is weak, but that there is a trend for those CS strategies with more memory slots to have lower cliquishness, lower weighting of the past, and higher provokability.

3.3. Cooperative strategies invading pure defectors (simulation set 2)

Simulation set 1 indicates that a cooperative strategy ($D = 0$) characterized by low τ , moderate δ , and high χ is successful in a wide range of environments. Based on this we examine whether ‘cliquer’-type strategies (with $\tau = 0.15$, $\delta = 0.50$, and $\chi = 1$) can invade a population of pure defectors when initially rare. For comparison, we also examine the ability of a fixed OFT-type strategy (with $\tau = 0.50$, $\delta = 0$, and $\chi = 0$) to invade a population of pure defectors when initially rare. We consider invasion under a set of baseline conditions ($N = 400$, $B/C = 2$ or 4 , $e_s = 0$ and 0.05 , $e_m = 0.05$), allowing the initial proportion of cooperative strategies to vary between 0.005 (2CS against 398 pure defectors) and 0.30 (120CS against 280 pure defectors). Each scenario was run 10 times.

In these invasion scenarios, reproduction occurred *without mutation*, and we counted that a CS successfully invaded when it completely displaced the population of defectors. When the population of defectors extinguished the CS, we counted that CS failed to invade.

In some OFT invasions the population reached a stable mix of strategies which we counted as a draw, recording the mean fraction of each strategy type.

In the four environments considered ($B/C = 2$ or 4 , $e_s = 0$ or 0.05 where $N = 400$ and $e_m = 0.05$) cliquer strategies consistently invaded and displaced a population of pure defectors with invading populations starting as low as 8 strategies out of 400 (Table 3). With smaller initial populations, decreasing B/C ratio and increasing strategy error decreased the probability of invasion and displacement by cliquers.

When errors were set to zero (i.e. when $e_m = 0$ and $e_s = 0$), Out-for-tat strategies invaded and replaced a population of unconditional defectors in 10 out of 10 runs with initial population prevalence as low as 0.01 (not shown in Table 3). However, adding either kind of error greatly reduced out-for-tat’s ability to invade. In environments with some strategy execution error ($e_s = 0.05$), out-for-tat was extinguished by unconditional defectors. In environments with some matching error ($e_m = 0.05$, but $e_s = 0$), out-for-tat strategies were able to increase their prevalence and reach a stable proportion of the population (mean proportion = 0.42 when $B/C = 2$ and 0.56 when $B/C = 4$), but they were unable to completely displace unconditional defectors.

3.4. Part-time defectors against cliquers (simulations sets 3 & 4)

Although cliquers are successful in emerging and persisting against classical strategies in these models, it is possible to imagine alternative strategies that might be able to infiltrate a population of cliquers. For example, a strategy that defects occasionally with some probability may well displace a population of cliquer strategies. This, in turn, may initiate a downward spiral of

increasing defection towards a population comprised entirely of unconditional defectors. To examine how such a *part-time defector* might initially invade a population of cliquers, we designed a strategy of similar complexity to that of the cliquer strategy and compared the selective advantage of cliquers and part-time defectors in the same way that we initially compared cliquers and pure defectors. Like cliquers, these part-time defectors use a weight of past experience (δ), a threshold for acceptable cooperation (τ) and a memory for up to $K = 3$ partners to identify and to cultivate cooperative partners. Unlike cliquers, they use χ to regulate their defection (rather than to regulate behavior in the random matching pool), deciding not to cooperate with a probability χ . We pitted these part-time defectors whose parameter values (δ , τ , χ) were permitted to evolve against a population of cliquers with fixed values close to those which had evolved in earlier simulations ($\tau = 0.15$, $\delta = 0.5$ and $\chi = 1$)—this gives a substantial edge to part-time defectors.

Under the range of conditions we examined ($B/C = 1$, 2 or 4, $e_s = 0.05, 0.10, 0.20$, $N = 100, 400$), cliquers had a selective advantage over the part-time defectors (maintaining proportions of 86–97% in the population), except when $B/C = 1$. However, the selective advantage of cliquer over part-time defectors was slightly weaker in environments with lower B/C ratios and higher rates of strategy error, results that mirror our findings when comparing cooperative strategies against unconditional defectors. Under the same range of environmental conditions, populations of part-time defectors (starting at 10% and 30%) were also unable to invade cliquer populations when mutation was turned off.

3.5. Distribution of relationship durations

In contrast to many evolutionary IPD models, our partner selection approach does not fix the probability (ω) that a relationship between two partners will continue to the next round (Axelrod and Hamilton, 1981; Boyd and Richerson, 1992; Brown et al., 1982). Rather, in our model interactions end ‘naturally’ as the result of changing partner preferences, the death of partners, and the end of the simulation. In addition, mutually preferred interactions may be temporarily interrupted by interactions with other partners. From these ‘naturally’ occurring relationships, it is possible to derive the actual distribution of relationship durations and to compare it to the geometric distribution assumed in most IPD models.

To illustrate, consider a situation where cliquers invade defectors ($e_s = 0.05$, $e_m = 0.05$, $B/C = 2$, $N = 100$). Over 200 generations (20,000 iterations), there were 49,038 relationships where two partners interacted for more than one round. Of these, 5298 ended after two interactions, while another 3780 ended




Fig. 6. Probability of relationship continuation (w) during first 40 interactions between two partners ($N = 100$, $e_s = 0.05$, $e_m = 0.15$, 200 generations).

after three interactions. From these dissolutions, we can calculate the probability of ending a relationship after a given number of interactions. Fig. 6 shows the change in ω during the first 40 interactions between two individuals. Although ω settles around 0.96 after the first 10 interactions, it shows a divergence from the assumption of uniform ω in early interactions, with higher rates of dissolution at early stages of a relationship.

4. Discussion

In this discussion section we (1) summarize our key findings, (2) review these findings vis à vis related work on the evolution of cooperation, (3) discuss the limitations of this work and lay out ongoing inquiries, (4) consider the implications of our results for interpreting recent experimental work on human sociality from behavioral economics, and (5) emphasize the importance of theorizing about relationships and relationship strategies, rather than individual traits, like ‘altruism’.

4.1. Summary of key findings

These results build on past work on the evolution of dyadic cooperation by studying an explicit partner choice model that (1) permits individuals to play the PD game with preferred partners, (2) parameterizes the strategy space (rather than limiting the process to a finite number of discrete strategies), (3) allows mutation on all evolving parameters permitting a full range of strategic ecologies,¹² (4) includes two kinds of behavioral noise, (5) limits partner memory, and (6) considers the effects of population size and drift. Our findings suggest that a cliquish strategy which is forgiving towards preferred partners but tends to defect in interactions with ‘strangers’ (non-partners) can invade and resist invasion by pure defectors in a range of environmental conditions. We also present tentative findings that the cliquish strategy can also resist invasion by more sophisticated part-time defectors.

¹²The ecology of possible strategies has been shown to affect drastically the success of particular strategies.

Broadly, our work confirms the most general finding from the canonical model: benefit-to-cost ratio show a steeply nonlinear threshold-like effect on the emergence of cooperation. However, in contrast to the canonical findings, we also demonstrate that certain kinds of cooperative strategies *can invade when rare and remain stable* even in the face of substantial behavioral noise. The difficulty of invasion that plagued early work appears to be due in part to the inability of individuals to increase their probability of playing the PD game with preferred partners (a restriction we feel is largely unrealistic for humans). In our model, rare *cliquers* can continue searching for other sufficiently cooperative partners and once found, they can potentially interact for as long as both partners (1) maintain mutually acceptable cooperation and (2) survive into the next round.

Successful *cliquer* strategies differ from TFT-like strategies in two ways. First, cliquers integrate past behavior into partner preferences, with about half of the weight on the most recent interaction and about half on prior interactions—note, these strategies are *not* remembering more information about past interactions; they are merely weighting the construction of their preferences differently. Interestingly, weighting of past outcomes was considerably weaker in environments with no strategy noise. This is consistent with past findings that TFT-like strategies are most successful in environments where there is no threat of execution errors (Bendor, 1987).

The second difference between cliquers and TFT-like strategies is that cliquers are not generally cooperative with new interactants (i.e. not NICE). If their K partner memory slots are filled, they defect on non-partners most of the time. Cliquishness contributes to a strategy's selective advantage in two important ways. First, it shields individuals from the possibility of exploitation by unknown partners in the random matching pool. By selectively defecting with strangers, a cliquish strategy avoids some of the risks of cooperating. Second, cliquish individuals who find themselves in the random matching pool also exploit NICE individuals who unconditionally cooperate with strangers. In contrast, defectors in a world of mostly cliquers will find fewer NICE strategies in the random pool to exploit.

Overall, in our simulations a strategy's ability to invade and remain stable in the face of noise and a wide-ranging ecology of other strategies seems to rest on its ability to cultivate and to maintain a small set of local relationships with known individuals, while largely ignoring (or exploiting) the rest of the population. In contrast, NICE strategies, such as OFT, get exploited by unconditional defectors and cliquers (with their K slots full) in the random pool.

We think this work offers both an intuitively more plausible evolutionary model for our species than the

IPD, and yields results consistent with the empirical patterns of dyadic cooperation observed in human social relationships (Silk, 2003). Specifically, successful strategies eschew tit-for-tat accounting by tracking a level of preference in a partner that is slowly built up or eroded based on repeated interactions (due to a high δ value). Although this mode of accounting is slow to break with an old friend, it is sensitive to early interactions, which makes it an effective antidote to full-time defectors. As just noted, the strategy is not unconditionally NICE and does not typically cooperate with non-partners (due to a high χ value).

4.2. Partner choice and assortative interactions as avenues to cooperation

The success of the cliquer strategy is a specific case of the general observation that cooperative strategies must have a greater than random chance of interacting with other cooperative strategies in order to avoid exploitation by free riders (Eshel and Cavalli-Sforza, 1982; Wilson and Dugatkin, 1997). Many past efforts at modeling this non-random matching, without recourse to kinship or reputation, have either (1) tacitly assumed cooperative strategies are able to identify and interact with other cooperative strategies (Eshel and Cavalli-Sforza, 1982; Peck, 1993; Wilson and Dugatkin, 1997) or (2) explicitly modeled the assortment process using spatial patterning (Eshel et al., 2000; Sella and Lachmann, 2000; Watanabe and Yamagishi, 1999) or tribal or dialectical assortment (Nettle and Dunbar, 1997). Our approach is different from these because it relies on the processes and resulting psychology of focusing effort on long-term enduring cooperative relations with *specific* partners, rather than determining whether to cooperate with a randomly chosen individual with whom one may (1) have several repeated interactions or (2) share a common attribute. We argue that this specification of partner choice is resistant to the roving defectors that might exploit spatially configured populations (Dugatkin and Wilson, 1991; Enquist and Leimar, 1993; Watanabe and Yamagishi, 1999) and the mimics that can exploit tribal or dialectical matching (McElreath et al., 2003; Nettle and Dunbar, 1997) because it is more difficult to mimic a specific person than to rove widely or to imitate a dialect or tribal identity.

Our approach to modeling partner choice also appears to resolve conflicting findings about the role that partner switching plays in the evolution of cooperation. In some situations, partner choice permits the avoidance of free riders and thus favors cooperation (Noe and Hammerstein, 1994). On the other hand, free riders can exploit partner switching by moving from partner to partner without fear of adverse consequences (Dugatkin and Wilson, 1991). Our model allows

cooperative strategies to switch partners at high rates until they find an acceptably cooperative partner, at which point their rate of switching decreases dramatically. This provides the flexibility needed initially to exclude free riders, and the stability necessary to avoid them in future interactions. Conversely, once cooperative strategies have cultivated a set of stable relationships (a ‘local environment’), defectors continue to switch between partners, but these partners are also more likely to be defectors since they too appear in the random matching pool. Consequently, the advantage to defectors given by partner switching is substantially diminished.

4.3. Limitations and future directions

Solving the dilemma of dyadic cooperation involves two interrelated trade-offs. First, an actor must decide to cooperate and face potential exploitation, or avoid cooperation altogether (Bendor, 1993). However, the possibility of cultivating a relationship introduces a new tradeoff—that between the increased certainties gained through building a few relationships and the potential opportunities lost by remaining confined to specific partners. In environments with heterogeneous payoffs, if one does decide to cooperate and to cultivate a relationship, it is necessary to monitor whether one is getting an adequate payoff from a current partner relative to other *potential* partners. This paper has focused on the first trade-off by allowing individuals to test and cultivate relationships. However, as social exchange theories have suggested, it is also important to understand how individuals make guesses about the opportunities they lose by engaging in one relationship versus another (Kelley and Thibaut, 1978). In the future, it will be important to model how individuals survey their environment and make decisions to move to new relationships based on their estimations.

Although cliquers are successful in emerging and persisting against traditional strategies, it is possible to imagine alternative strategies that might be able to infiltrate a population of cliquers. In this paper, we examined how one alternative strategy, a part-time defector, fares against cliquers. Of course, there are likely numerous ways in which one could implement a strategy like part-time defector of comparable complexity to cliquer, and some of these strategies may be more successful than the one we proposed here. However, while it is possible to imagine part-time defectors that could displace a population of cliquers, it is equally possible to describe adaptations that might allow cliquers to resist such an invasion. For example, cliquers may persist by adjusting their long-term memory weightings (δ) and cooperative thresholds (τ) to counter-balance these part-time defectors. In a cultural species like humans, cliquers also may use a second system of

cultural inheritance to rapidly adjust their parameters (δ , τ) to the latest incarnation or distribution of part-time defectors (Henrich and Henrich, in press). Ultimately, finding part-time defectors which are able to invade and displace populations of cliquers would give us a greater understanding of the conditions under which cooperation thrives and under which it can be exploited by defecting strategies of comparable complexity. Future simulation work will examine how cooperative strategies might adapt to such part-time defectors with evolving values for δ and τ or a second system of inheritance.

4.4. Does reciprocal altruism explain prosocial behavior in one-shot economic experiments?

Findings from this model bear on a current debate at the interstices of evolutionary biology and behavioral economics. Results from experiments in behavioral economics showing prosocial behavior in anonymous one-shot experimental games, such as the ultimatum game, prisoner’s dilemma, and public goods games, have been interpreted as inconsistent with the predictions of reciprocal altruism and in need of an alternative evolutionary explanation (Fehr and Fischbacher, 2003). Critics of this have countered by suggesting that reciprocal altruism can explain people’s behavior once one accounts for the fact that human social psychology evolved in smaller-scale societies with stable long-term membership. In these environments, some have reasoned (presumably from the results of the canonical IPD models) that individuals should have evolved to be NICE (i.e. always cooperate in ‘round 1’ of a repeated game) to everyone, because anyone they encounter is likely to be a long-term future interactant. Carrying this evolved psychology to modern social environments and behavioral experiments, the argument suggests that people cooperate in one-shot games (are NICE) because they are acting ‘as if’ they are in a repeated sequence of interactions—even though they are not (Johnson et al., 2003).

Testing the theoretical logic underlying this argument, our model—which is based on the logic of repeated interaction—predicts that people should *not* typically cooperate with just anyone from their group (i.e. they should not generally be NICE to anyone). If the above theoretical proposal were sound, χ in our model would tend to stabilize near a value of zero. Instead, χ remains well above 0.5, even for extremely high values of B/C . In fact, we did not find any conditions favorable to the $\chi = 0$ supposition. Given this, we submit that the claim of NICENESS is an artifact of the independence assumption of interactions in the IPD. Thus, the behavior in one-shot experiments cannot likely be explained by the

evolutionary logic of repeated interaction on which reciprocal altruism is founded¹³.

4.5. Thinking about relationships

Theory in the evolution of cooperation has been influenced by the notion that ‘altruism’, ‘trustworthiness’, ‘reciprocity’ and ‘reputation’ are properties of the individual. However, in daily life we frequently observe individuals extend a helping hand to specific companions, while neglecting to aid others. Some individuals may have a bad reputation in general, but still be a reliable friend and cooperator with a select few. Our model begins to capture this by producing individuals that are ‘cooperative’ and ‘altruistic’ with some people, but not with most. In a population of cliquers, for example, no cliquer is more or less altruistic than the others, yet there is a great deal of heterogeneity in who is altruistic to whom. The possibility for altruistic interactions exists precisely because individuals can cultivate *cooperative relationships*. This is not to say that there are not trait-like differences between individuals, but rather to shift the emphasis to the relationship-building aspects of our species. Further exploration of how individuals cultivate cooperative local environments would be a useful addition to the current emphasis on individual traits.

Acknowledgments

We would like to thank Karthik Panchanathan, Rob Boyd, Kenneth Maes and participants of the Evolution and Human Behavior seminar at Emory University for helpful comments. We are also grateful for the careful and constructive comments made by two anonymous referees. DJH acknowledges financial support from an NSF dissertation improvement grant.

References

- Ashlock, D., Smucker, M.D., Stanley, A.E., Tesfatsion, L., 1996. Preferential partner selection in an evolutionary study of Prisoner's dilemma. *Biosystems* 37, 99–125.
- Axelrod, R., 1984. *The Evolution of Cooperation*. Basic Books, New York.
- Axelrod, R., Hamilton, W., 1981. The evolution of cooperation. *Science* 211, 1390–1396.
- Batali, J., Kitcher, P., 1995. Evolution of altruism in optimal and compulsory games. *J. Theor. Biol.* 175, 161–171.
- Bendor, J., 1987. In good times and bad: reciprocity in an uncertain world. *Amer. J. Polit. Sci.* 31, 531–558.
- Bendor, J., Swistak, P., 1997. The evolutionary stability of cooperation. *Amer. Polit. Sci. Rev.* 91, 290–307.
- Bendor, J., Kramer, R., Swistak, P., 1996. Cooperation under uncertainty: what is new, what is true, and what is important. *Am. Sociol. Rev.* 61, 333–338.
- Boerlijst, M.C., Nowak, M.A., Sigmund, K., 1997. The logic of contrition. *J. Theor. Biol.* 185, 281–293.
- Boyd, R., Lorberbaum, J.P., 1987. No pure strategy is evolutionarily stable in the repeated Prisoner's dilemma game. *Nature* 32 (6117), 58–59. *Nature*, 32:58–59.
- Boyd, R., Richerson, P., 1992. Title: Punishment allows the evolution of cooperation (or anything else) in sizable groups. *Ethol. Sociobiol.* 13, 171–195.
- Brown, J.S., Sanderson, M.J., Michod, R.E., 1982. Evolution of social behavior by reciprocity. *J. Theor. Biol.* 99, 319–339.
- Bshary, R., Noe, R., 2003. Biological markets: the ubiquitous influence of partner choice on the dynamics of cleaner-fish-client fish interactions. In: Hammerstein, P. (Ed.), *Genetic and Cultural Evolution of Cooperation*. MIT Press, Cambridge, MA, pp. 167–184.
- Connor, R.C., 1992. Egg-trading in simultaneous hermaphrodites: an alternative to tit-for-tat. *J. Evol. Biol.* 5, 523–528.
- de Vos, H., Smaniott, R., Elsas, D.A., 2001. Reciprocal altruism under conditions of partner selection. *Ration. Society* 13, 139–183.
- de Vos, H., Zeggelink, E., 1997. Reciprocal altruism in human social evolution: the viability of reciprocal altruism with a preference for “old-helping-partners”. *Evol. Hum. Behav.* 18, 261–278.
- Dugatkin, L.A., Wilson, D.S., 1991. Rover: a strategy for exploiting cooperators in a patchy environment. *Amer. Nat.* 138, 687–701.
- Dunbar, R.I.M., 1998. The social brain hypothesis. *Evol. Anthropol.* 6, 178–190.
- Enquist, M., Leimar, O., 1993. The evolution of cooperation in mobile organisms. *Anim. Behav.* 45, 747–757.
- Eshel, I., Cavalli-Sforza, L.L., 1982. Assortment of encounters and evolution of cooperativeness. *Proc. Natl Acad. Sci. USA.* 79, 1331–1335.
- Eshel, I., Herreiner, D.K., Samuelson, L., Sansone, E., Shaked, A., 2000. Cooperation, mimesis, and local interaction. *Sociolog. Methods Res.* 28, 341–364.
- Fehr, E., Fischbacher, U., 2003. The nature of human altruism. *Nature* 425, 785–791.
- Fehr, E., Henrich, J., 2003. Is strong reciprocity a maladaptation? In: Hammerstein, P. (Ed.), *Genetic and Cultural Evolution of Cooperation*. MIT Press, Cambridge.
- Feldman, M., Thomas, A., 1987. Behavior-dependent contexts for repeated plays of the prisoner's dilemma. II. Dynamical aspects of the evolution of cooperation. *J. Theor. Biol.* 128, 297–315.
- Friedman, J.W., Hammerstein, P., 1991. To trade, or not to trade; that is the question. In: Selten, R. (Ed.), *Game Equilibrium Models I: Evolution and Game Dynamics*. Springer, Berlin, pp. 257–275.
- Getty, T., 1987. Dear enemies and the prisoner's dilemma: why should territorial neighbors form defensive coalitions. *Amer. Zoologist* 27, 327–336.
- Gotts, N.M., Polhill, J.G., Law, A.N.R., 2003. Agent-based simulation in the study of social dilemmas. *Artif. Intell. Rev.* 19, 3–92.
- Hammerstein, P., 2003. Why is reciprocity so rare in social animals. In: Hammerstein, P. (Ed.), *Genetic and Cultural Evolution of Cooperation*. MIT Press, Cambridge.
- Hauk, E., 2001. Leaving the prison: permitting partner choice and refusal in prisoner's dilemma games. *Comput. Econ.* 18, 65–87.
- Hayashi, N., Yamagishi, T., 1998. Selective play: choosing partners in an uncertain world. *Personality Social Psychol. Rev.* 2, 276–289.
- Henrich, N.S., and J. Henrich, in press. *The Cultural and Evolutionary Origins of Cooperation: Kinship, Reciprocity, Social Norms, and Ethnicity among the Chaldeans of metro Detroit*. Oxford University Press, Oxford.
- Johnson, D., Stopka, P., Knights, S., 2003. The puzzle of human cooperation. *Nature* 421, 911–912.

¹³There are other empirical and theoretical problems with using reciprocal altruism to explain the prosocial behavior observed in one-shot games (Fehr and Henrich, 2003)

- 1 Kelley, H.H., Thibaut, J.W., 1978. *Interpersonal Relations: A Theory of Interdependence*. Wiley, New York.
- 3 Leimar, O., Hammerstein, P., 2001. Evolution of cooperation through indirect reciprocity. *Proc.: Biol. Sci.* 268, 745–753.
- 5 Lorberbaum, J., 1994. No strategy is evolutionarily stable in the repeated prisoner's dilemma. *J. Theor. Biol.* 168, 117–130.
- 7 Lorberbaum, J.P., Bohning, D.E., Shastri, A., Sine, L., 2002. Are there really no evolutionary stable strategies in the iterated prisoner's dilemma. *J. Theor. Biol.* 214, 155–169.
- 9 Markovits, H., Benenson, J.F., Kramer, D.L., 2003. Children and adolescents' internal models of food-sharing behavior include complex evaluations of contextual factors. *Child Dev.* 74, 1697–1708.
- 11 McElreath, R., Boyd, R., Richerson, P.J., 2003. Shared norms and the evolution of ethnic markers. *Curr. Anthropol.* 44, 122–129.
- 13 Milardo, R.M., 1992. Comparative methods for delineating social networks. *Social and Personal Relationships* 9, 447–461.
- 15 Nettle, D., Dunbar, R., 1997. Social markers and the evolution of reciprocal exchange. *Curr. Anthropol.* 38, 93–99.
- 17 Noe, R., Hammerstein, P., 1994. Biological markets: supply and demand determine the effect of partner choice in cooperation, mutualism, and mating. *Behav. Ecol. Sociobiol.* 35, 1–11.
- 19 Peck, J.R., 1993. Friendship and the evolution of cooperation. *J. Theor. Biol.* 162, 195–228.
- 21 Sella, G., Lachmann, M., 2000. On the dynamic persistence of cooperation: how lower individual fitness induces higher survival. *J. Theor. Biol.* 206, 465–485.
- 23 Sherratt, T.N., Roberts, G., 1998. The evolution of generosity and choosiness in cooperative exchanges. *J. Theor. Biol.* 193, 167–177.
- 31 Silk, J., 2003. Cooperation without counting: the puzzle of friendship. In: Hammerstein, P. (Ed.), *Genetic and Cultural Evolution of Cooperation*. MIT Press, Cambridge, MA, pp. 37–54.
- 33 Skyrms, B., Pemantle, R., 2000. A dynamic model of social network formation. *Proc. Natl Acad. Sci. USA* 97, 9340–9346.
- 35 Stanley, A., Ashlock, D., Tesfatsion, L., 1994. Iterated prisoner's dilemma with choice and refusal of partners. In: Langton, G. (Ed.), *Artificial Life III*. Addison-Wesley, New York.
- 37 Sugawara, K., 1984. Spatial proximity and bodily contact among the Central Kalahari San. *African Study Monographs* 3, 1–43.
- 39 Trivers, R.L., 1971. The evolution of reciprocal altruism. *Q. Rev. Biol.* 46, 34–57.
- 41 Watanabe, Y., Yamagishi, T., 1999. Emergence of strategies in a selective play environment with geographic mobility: a computer simulation. In: Foddy, M. (Ed.), *Resolving Social Dilemmas: Dynamics, Structural, and Intergroup Aspects*. Edwards Brothers, Ann Arbor, MI, pp. 55–66.
- 43 Wilson, D.S., Dugatkin, L.A., 1997. Group selection and assortative interaction. *Amer. Nat.* 149, 336–351.
- 45 Yamagishi, T., Hayashi, N., Jin, N., 1994. Prisoner's dilemma networks: selection strategy versus action strategy. In: *Social dilemmas and cooperation*. Springer, New York, pp. 233–250.
- 47
- 49

UNCORRECTED PROOF