Journal of Economic Behavior & Organization
Vol. 44 (2001) 145–167

JOURNAL OF
Economic Behavior
& Organization

www.elsevier.com/locate/econbase

Help, harm or avoid? On the personal advantage of dispositions to cooperate and punish in multilateral PD games with exit

Roger D. Congleton*, Viktor J. Vanberg[1]

Center for Study of Public Choice, George Mason University, Fairfax, VA 22030, USA

Received 19 October 1998; received in revised form 19 October 1999; accepted 26 October 1999

Abstract

Simulated round-robin tournaments are used (1) to assess the relative performance of non-defecting, conditionally cooperative programs in a three-person PD setting where exit is possible (PDEs) and (2) to assess the importance of sanctions for the success of joint enterprises. We find that the possibility of exit allows individuals to escape from dysfunctional teams which increases the potential benefits of cooperation while reducing those associated with free-riding. We also find that punishment, especially targeted punishment, is critical to the success of shorter term enterprises. Similar results held in evolutionary settings where successful behavioral programs can be imitated, or propagated by non-rational biological or social processes. Our results suggest that non-defecting programs of conditional cooperation can be rational, average payoff maximizing, strategies for participating in small multilateral settings where individuals can choose to participate in a PD game or not. © 2001 Elsevier Science B.V. All rights reserved.

JEL classification: C7; D6; A13

Keywords: Cooperation; Evolution; Exit; Ethics; Economics; Team production; Public goods; Punishment; Sanctions; Simulations; Game theory; Clubs

---

1. Introduction

Experimental evidence from prisoner's-dilemma games indicates that many, but not all, players cooperate in both repeated and non-repeated PD settings. This very robust

---

* Corresponding author.

[1] Center for Study of Public Choice, George Mason University, Fairfax, VA 22030, USA; and Institut für Allgemeine Wirtschaftsforschung, Universitat Freiburg, Freiburg i.Br. Germany. The authors have benefited from discussions and comments made at several academic meetings and seminars and the very helpful criticisms of two referees.

0167-2681/01/$ – see front matter
PII: S0167-2681(00)00159-1

experimental result is a puzzle for rational choice theorists. Incentives affect cooperation in a manner consistent with economic intuitions, so it is clear that rational decision making takes place at the margin, but "too much" cooperation takes place for the predictions of elementary non-cooperative game theory, while too little takes place for some more sophisticated dynamic theories.[2] One possible explanation of the relatively extensive cooperation observed in experimental settings is that players use the same behavioral rules when participating in PD experiments that they use in their ordinary lives. Such rules may be broadly rational in the sense that the outcomes achieved are better for the individual, *on average*, than the results of local optimization. It is possible that such personal behavioral routines, or dispositions, may *mistakenly* call for cooperation in experimental PD settings. On the other hand, it bears noting that individual and group payoffs are, of course, *much higher* in the experiments than would have been achieved with the non-cooperative behavior that "rationality" recommends in finite PD settings.[3]

The "ruleful action" interpretation is consistent with that computer simulation experiments evaluate the relative performance of alternative programs for playing two-person PD games. An implication of Axelrod (1984) finding regarding the relatively superior performance of cooperation oriented "Tit-for-Tat" and "Tit-for-two-Tat" programs is that individuals who rely upon programs of conditional cooperation will do better in life's protracted two-person PD tournaments than those who uniformly defect. Bendor et al. (1991) note that relatively forgiving cooperative programs may perform better in circumstances where PD payoffs are themselves uncertain. Vanberg and Congleton (1992) indicate that the possibility of exit makes programs of conditional cooperation more viable as strategies for playing ordinary PD games.[4]

On the other hand, it may be argued that the simulation results of Axelrod and the others are dependent on the two-person games analyzed and thus provide little insight into the evolution

---

[2] In repeated game settings, whether cooperation can be a rational strategy depends largely on expectations of reciprocity or sanctions. Textbook explanations include those in Kreps (1990, Chapter 14) and in Mas-Colell et al. (1995, Chapter 9). Explanations of cooperative behavior in PD games where 'types of players' exist have also been developed in theoretical work. See, for example: Kreps et al. (1982). A nice overview of the relevance of Nash equilibrium from the perspective of evolutionary game theory is developed by Mailath (1998).

[3] Several recent experimental studies have attempted to explain the excessive cooperation observed in PD experiments. Andreoni and Miller (1993) examine the sequential equilibrium reputation hypothesis in the finitely repeated prisoner's-dilemma. Their results suggest an important role for what they call "homemade altruism," that is to say, of preexisting predispositions to cooperate. Cooper et al. (1996) investigate cooperative play in prisoner's-dilemma games. They argue that the observed level of cooperation can be explained as a consequence of reputation building and altruism. Nalbantian and Schotter (1997) examine how cooperation in joint enterprises is affected by compensation schemes. Our study explores whether the observed propensities to cooperate can, in fact, be *rational* strategies for participating in small repeated multilateral PDE games.

[4] The spirit of our analysis is consistent with Simon (1979) various studies of behavior within firms which suggests that rather than "optimize" employees use various "heuristics" to get their jobs done, and with Hayek (1978) discussion of the cultural counter part to genetic evolution. Simon's argued that heuristics tend to be used because the problems at hand are too complex for individuals to solve them rationally. Hayek suggests that durable rules of conduct must have performed well in a variety of circumstance in order to survive from one generation to the next. Our analysis suggests that individuals employ heuristics, rules of thumb, behavioral routines, and so forth in settings of joint production because they can achieve better results on average with those decision rules than would have been achieved with case by case, *but local*, optimization. See also Rubin (1982), Hirshleifer and Coll (1991) and Sugden (1986).

of real world behavioral programs.[5] There is an inherent reciprocity in two-person games that is lacking in games between larger groups of players. As the number of players increases, it becomes increasingly difficult to observe and sanction non-cooperative behavior, and it may be that these two possibilities account for much of the success of programs of conditional cooperation in various two-person simulation tournaments. Critics may argue that the human and simulation experimental results from two-person games have limited relevance for settings where more than two players interact.

To explore the possibility that the relative performance of programs of conditional cooperation is so limited, this paper analyzes the private merits of cooperative behavior in small *multilateral* prisoner's-dilemma (PD) situations in market-like settings where *exit* is possible. Multilateral prisoner's-dilemma games with exit (PDE games) are exemplified, for instance, by the team production and collective goods problems that private firms, cooperatives, and clubs must overcome to be successful joint enterprises. Such joint activities yield a team surplus if a sufficient number of team members cooperate, but yield less than what the participants could realize in separate action if all free-ride or defect.[6] As in any PD setting, the result when all participants contribute is clearly preferable for all involved to a situation where all seek to free-ride, yet, the PD nature of incentives works against the cooperative outcome. The principle question investigated below is the extent to which programs of conditional cooperation become less successful in multilateral PDE games insofar as detecting and sanctioning specific free-riders in a repeated game setting is significantly more difficult.

The present analysis parallels our earlier study in many respects, Vanberg and Congleton (1992). We again use an Axelrod-type simulation tournament to determine the viability of alternative behavioral programs. The viability or rationality of alternative behavioral programs is again judged, not from deductive properties, but from relative performance in simulated PDE tournaments. We again find evidence of the importance of exit in facilitating cooperative solutions to PDE games. What is new is the extension of our analysis to multilateral games where many of our arguments concerning reciprocity, punishment, and signaling are clearly weaker, and the potential complexity of interactions among strategies is clearly much greater. We also examine the merits of a new retributive behavioral program that punishes only free-riders, but at a personal cost. Such programs become noteworthy only in a multilateral setting, and may, as developed below, be a personally profitable method of increasing cooperation within a multilateral PDE setting.

---

[5] As Axelrod (1984: 11) notes about his "The Evolution of Cooperation": "In this book I will examine interactions between just two players at a time. A single player may be interacting with many players, but the player is assumed to be interacting with them one at time." The relevant contrast between multilateral and bilateral PDs is not simply the number of players involved. Whether or not a multi-person setting can be disaggregated into bilateral relations depends on the nature of the underlying "problem of cooperation."

[6] The notion of a team surplus can be interpreted in terms of A. Alchian's and H. Demsetz's (1972) theory of team production. In their study of multi-player PDs Hirshleifer and Rasmusen (1989) speak of "advantage to grouping" and of "aggregation economies." It bears noting that not every multi-person game with PD like payoffs is a multilateral PD game in the sense used in the present paper. Many games that at first glance appear to be multilateral can, on closer examination, be "factored down" into a series of bilateral PD games without significant analytical loss. This is true, for example, of many trading networks in ordinary markets.

The analysis is organized as follows: Section 2 describes the joint enterprise setting and the behavioral programs used for the round-robin tournaments. Section 3 presents the results of static tournament simulation experiments for a range of more or less durable joint enterprises.[7] Section 4 presents the results of evolutionary simulations for a setting where the composition of the population of behavioral programs reflects the relative performance of the programs in past tournaments/generations. Section 5 provides a summary of the results and discusses some possible applications.

The main results of the present paper are taken from a series of simulation experiments that assess the relative performance of nine behavioral programs for participating in three-person joint enterprises. The relative performance of the behavioral programs is first assessed using average scores in round-robin tournaments of three-person teams in a setting where teams may engage in joint production. As we will see, whether free-riding or cooperation is the "best strategy" varies with the nature of the joint production process, the composition of the teams find and the potential duration of the joint enterprise. No single behavioral program is found to dominate every possible setting for joint production, although several programs of conditional cooperation perform very well in a variety of settings. Generally, programs of conditional cooperation dominate other possible routines for participating in joint enterprises when those enterprises last for more than a few periods and participation in the enterprise is entirely voluntary.

2. Simulation parameters: the joint enterprise and dispositions to cooperate

2.1. The joint enterprise

The game analyzed in this paper is that associated with joint production in a setting where individual contributions are difficult to assess. Both the Buchanan (1965) theory of clubs and the Alchian and Demsetz (1972) characterization of team production indicate that various economies of scale can motivate voluntary production in groups.[8] Economies of group production may arise for a variety of reasons: the shared use of fixed capital goods, increases in specialization made possible by scale, organizational economies, and so forth. For concreteness, we assume that 'jointness' of production arises because the production process of interest requires some fixed factor(s) of production. As a joint enterprise, team members share the cost of the fixed factor(s) and, also, provide variable inputs to the joint enterprise. Whether members contribute or not to the shared fixed cost of the enterprise is assumed to be far easier to discern than the extent to which members provide variable inputs. The

---

[7] The simulations were run using a program written in Basic for the purposes of this paper. Its main task is keeping track of possible team combinations, computing player scores, compiling data, and computing population shares based on the tournament results (for the evolutionary simulations reported in Section 4). Behavioral programs are subroutines so that alternative strategies can be readily added to the tournament. Copies of the program are available upon request.

[8] Buchanan (1965: 1) characterizes clubs as "consumption ownership-membership arrangements" in which goods are "available for consumption to the whole membership unit" (ibid.: 3), and he notes that firms can also be considered as "one form of club organization" (ibid.: 5n). A. Alchian (1987: 1032): "The 'firm' . . . is a contractually related collection of resources of various cooperating owners. Its distinctive source of enhanced productivity is 'team' production, wherein the product is . . . a non-decomposable, non-attributable value produced by the group."

team's output is assumed to be shared among all team members as is often the apportioning rule used within partnerships, cooperatives, and other small voluntary organizations.

We use a particular concrete joint production function that can be parameterized to generate payoffs comparable to those used in previous simulation studies of PD and PDE games. The total output of a particular team in a particular production session is assumed to be twice the level of variable inputs provided to that session, $Q = 2\sum L_i$, where L_i can be thought of as a difficult to observe input such as labor intensity or quality. (The fixed factor allows production to take place but does not directly affect the marginal product of the variable input. Examples of such fixed factors include such inputs as trails, dikes, freight elevators, office space, printers, or web hosting services, all of which require regular maintenance and/or rental payments.)

Each team member's net payoff, P_i, is one third of the team's output net of fixed costs, F, less the opportunity cost, C_i, of contributing units of the variable input L_i.

$$P_i = \left(\frac{2\sum L_i - F}{3}\right) - C_i L_i \tag{1}$$

Note that if individual i free-rides, his variable costs fall by more than his claim to the firm's net output whenever $C_i > 2/3$. This creates a free-rider problem for the joint enterprise. Note also that the fixed-cost parameter of the payoff function can be manipulated to make the fruits of joint production more or less difficult to achieve. The greater is the fixed cost, the more difficult it is to realize the advantages of joint production.

Initially, it is assumed that fixed costs consume two units of output, $F = 2$, and that the opportunity cost of labor is one unit of output, $C_i = 1$. This implies that a three-person team is the smallest that can be mutually advantageous. A two-person team would not realize any advantages from joint production since: $(2 \times 2 - 2)/2 - 1 = 0$. Players first decide whether to participate or not in a particular production session (to contribute toward the fixed-capital good), and then decide whether to provide labor to the joint enterprise if the capital good is acquired. Purchase of the capital good is assumed to entitle the contributor to a share of the group's output. The non-participation, *or exit*, of a single player prevents joint production from taking place during the session of interest by eliminating the potential advantage of acquiring the fixed or shared factor of production. If all "team members" decide to participate in a particular joint-production session, then each team member independently makes a decision regarding his effort level, that is, to cooperate or free-ride. To *cooperate* in the joint enterprise means that the team member contributes 1 unit of labor to the production session. To free-ride means that the team member contributes 0 units labor to the joint production session. Free-riding by one or more parties does not rule out joint production since the fixed capital has already been purchased. Free-riding can only be detected *ex post*, after a production session has been completed.

A particular session of joint production may be judged a failure from the perspective of an individual team member if his net payoff is below zero, the assumed result without participation. A particular joint enterprise or team may be said to end if an exiting player never again contributes to the fixed factor, e.g. never returns.

Table 1 characterizes the various payoffs to contributors and free-riders on a three-person production team in the case where fixed costs consume two units of output. The magnitude of the payoffs is comparable to that used in Axelrod's two-person tournaments. Note that this

Table 1

Combination of strategies	Payoff for Cooperator	Payoff for Free-Rider
All Cooperators	0.33	–
Two Cooperators and one Free-Rider	−0.33	0.667
One Cooperator and two Free-Riders	−1	0
All Free-Ride	–	−0.66
One or more Exit	0	0

joint enterprise game resembles a prisoner's-dilemma game in the absence of an exit option insofar as free-riding is the dominant strategy ignoring the exit option. A move from any of the left-hand cells to the right-hand one below increases a player's single round payoff. The possibility of exit eliminates the strict dominance of free-riding, although free-riding continues to dominate cooperation in the case where two other players provide labor to the joint enterprise ($0.66 > 0.33$), and free-riding is as profitable as exit if one other team member free-rides. It is less profitable than exit if two others free-ride.

If one or more players *exit*, production does not take place, and all players receive a zero payoff until the exiting player(s) return to the team. Exit, consequently, secures a payoff that is better than that obtained from cooperating in a group which includes two defectors. However, exit is less rewarding than defection or cooperation on teams where the other two team members cooperate. Successful joint enterprises would not exist under the conventional rationality assumptions employed by economists and game theorists because mutual cooperation is not an equilibrium of this game.

It bears noting that this joint production process is a *very demanding setting* for joint enterprises. Only in the case where all members contribute to the joint production does the value of output exceed total cost, and only in that case do all members benefit from membership in the joint enterprise. A less demanding production setting is examined below in subsection IIIC.

2.2. Nine behavioral programs

The simulations evaluate the performance of nine behavioral programs for participating in the three-person joint enterprises characterized above. The behavioral programs included are essentially those used in Vanberg and Congleton (1992) modified for the multilateral PDE setting. The use of similar behavioral programs allows us to compare the results of two-person and three-person PDE tournaments without presenting a long series of parallel results. As we will see, many of the new results are surprisingly similar to those obtained for our previous two-person based study (Table 2).

The nine programs included vary in the sophistication with which they adapt to outcomes of previous rounds of play, or joint production sessions, and with respect to the range of strategies that they may employ. We regard the programs to be "ideal types" which may be readily generalized and understood by the reader. The latter is especially important in simulation studies where the interaction is complex and results need to be interpreted as much as computed. Transparency is one advantage of our approach over the genetic algorithm method where equilibrium strategies are often difficult to classify or to make sense of.

Table 2

Program	Strategy
Cooperator	Always cooperate. Always participate.
Free-Rider	Always defect. Always participate.
Tit-for-Tat	Cooperate in the first round, respond with defection when the payoff falls below zero, and with cooperation when the payoff is at least zero. Always participate.
Prudence	Always cooperate. Exit permanently from group if the current payoff is negative.
Hit-and-Run	Always defect. Exit permanently from group if the current payoff is negative.
Moral Suasion	Always Cooperate. Exit from group for one session if the current payoff is negative, then participate. Exit permanently if the payoff is or becomes negative after rejoining.
Easy-Tester	Defect initially, and thereafter if the payoffs are *positive*. Cooperate twice if the payoff in the last period was *not positive*; cooperate thereafter if the payoff is not negative, otherwise exit.
Recalcitrant-Tester	Defect initially, and thereafter if payoffs are *not negative*. Cooperate twice if the payoff in the last period *was negative*, cooperate thereafter if the payoff is positive, otherwise exit.
Targeted Retribution	Cooperate initially. Continue cooperating with other cooperators. Exit if there are two defectors. Punish a single defector. Exit if that defector does not cooperate in the following production session.

The two least sophisticated strategies included in the simulations are completely non-adaptive programs which either unconditionally defect (Free-Rider) or unconditionally cooperate (Cooperator). We use the terms free-riding and defection interchangeably. Unconditional defection, as noted above, is the “economically rational” strategy for the settings examined, *ignoring exit*. When exit is included, free-riding is an undominated strategy for two of the three possible choices that the other players may have made (all cooperate, one cooperate and one defect). However, the always free-ride strategy does not score as well as elementary game theory seems to suggest. Unconditional cooperation, is never an undominated strategy in this game, but in some settings it out-performs the unconditional defection program.

The other seven programs are all adaptive programs with more or less restricted ranges of strategic choices. All but one of the adaptive programs make strategy choices that are conditioned on their direct experience on a particular team. They are all informationally undemanding conditional strategies. Three of the adaptive programs make simple binary choices based entirely on their payoff in the previous production period. The first of these is “Tit-for-Tat,” the winner of the first Axelrod tournament. In a repeated two-person PDE game, Tit-for-Tat initially cooperates and then simply adopts the strategy previously used by his opponent. Tit-for-Tat, thus, either cooperates or defects. It never exits. In a multilateral PD game, Tit-for-Tat can be more or less tolerant in its response to defection because the other two players may not *all* be cooperating or free-riding.[9] We use a fairly tolerant version

---

[9] The notion of reciprocity is significantly different in two persons and multilateral PDE games. When two persons interact in a PD setting, their actions are necessarily targeted. Whatever one of them does, it affects, for good or ill, his particular counterpart. To respond to the other party’s choice by cooperating, means to encourage whatever it was the other party did. In such settings a reciprocating strategy like Tit-for-Tat can, as Axelrod’s tournament demonstrated, effectively generate cooperation among egoists in pairwise interactions. By contrast, in a multilateral PD it is not as clear what it means to reciprocate. Nor does a response send an unambiguous message to fellow team members in the multilateral PDE game.

of Tit-for-Tat in our tournaments. It initially cooperates, and defects only in response to a *negative* payoff in the previous production session.[10]

The other two elementary adaptive programs use the exit option. The cooperative program is called "Prudence." Prudence either cooperates or exits. The Prudence program cooperates whenever it participates in joint production sessions, but permanently exits from groups in which the previous payoff from cooperation was negative. The defection prone version of Prudence is called "Hit-and-Run." It either defects or exits. "Hit-and-Run" Free-Rides (defects) whenever it chooses to participate in joint production sessions, but permanently exits from groups in response to a negative payoff.

The other four adaptive programs use somewhat more complex strategies conditioned on more than past session payoffs. All have the explicit aim of changing the behavior of other group members. "Moral Suasion" is an extension of the Prudence program. It cooperates whenever it participates in production sessions, and it exits from a particular team when the payoff from cooperation in the previous session with that team was negative. However, rather than exiting forever, Moral Suasion returns to the group after exiting for one production session in order to determine whether fellow team members have responded to its exit by becoming cooperators. Moral Suasion is, thus, potentially able to profit from defectors who turn over a "new leaf" and contribute to the joint enterprise, although Moral Suasion runs the risk of returning to teams whose defecting team members have *not* changed their behavior.[11] If negative payoffs are realized after its return, Moral Suasion exits permanently from the team.

The defection-inclined counterpart to Moral Suasion is based on a program that we called Tester in our previous study. We use two versions of Tester in this tournament. "Easy-Tester" begins with defection and continues defecting as long as it receives *non-negative* payoffs. In response to a negative payoff Easy-Tester cooperates for two periods (in order to allow a Tit-for-Tat or Moral Suasion program to switch back into their cooperative mode) and then exits if sufficient cooperation is not elicited. "Recalcitrant-Tester" behaves similarly but switches away from the defection strategy only if scores are *negative*. Both strategies are teachable, in the sense that they can be converted from defection to cooperation by appropriate sanctions. Recalcitrant-Tester is a somewhat less responsive program than Easy-Tester. The Easy-Tester program was one of the highest scoring strategies in our

---

[10] Responding to payoff thresholds is equivalent to responding to particular numbers of defectors in any series of tournaments using a single payoff function. Thus, for a tournament based on a fixed payoff schedule it is of little importance what is chosen as the basis of response. Strategic responses to defection or payoffs do tend to vary across payoff functions. For example, consider the response of Tit-for-Tat. Tit-for-Tat responds by defecting if payoffs obtained in the previous round are negative. Two (or more) defectors would not elicit a defection from a payoff based Tit-for-Tat as long as the payoff associated with being a cooperator remains positive. If a single defection yields negative payoffs for the remaining cooperators, as in most of the settings explored here, then Tit-for-Tat strategy responds to a single defection with defection. A Tit-for-Tat program conditioned on number of defectors may have defected in the first case but not in the second. In either case, Tit-for-Tat programs may be more or less tolerant of defection.

[11] This aspect of strategies of conditional cooperation can be compared in terms of the protection that they provide against two kinds of error, analogous to the two types of errors in statistical testing. A type 1 error is to give up too early on a group that includes potential cooperators; a type 2 error is to continue too long with an unconvertible group of defectors. The success of behavioral programs can be expected to depend on how well they protect against both errors.

previous two-person simulations, and surprised us by being the strategy that *suffered most* from the shift to multilateral PDE games.

All of the above adaptive programs base their strategic adjustments on the past *payoffs* rather than on the past number of free-riders on the current team. We use past payoffs rather than number of defectors as the conditioning factor because past payoffs directly measure the private returns to participating on a given team and is the least informationally demanding index of team performance. Within a given joint production setting, behavioral programs conditioned on payoffs and ones conditioned on the number of defectors can be equivalent insofar as the *number* of defectors can always be deduced from the payoffs realized. However, it bears noting that this behavioral equivalence may not exist across joint enterprises insofar as payoffs from cooperation and defection vary as developed below in subsection IIIC. All the above programs also choose to cooperate or not *against the entire team* rather than against specific players as in a two-person game. They can not target their sanctions at particular team members.

The last adaptive program that we consider has the ability to target sanctions at particular players on their team. Such programs differ from those considered in our previous studies because the ability to target sanctions at specific players requires information about fellow team-member decisions that can not be deduced from a player's own payoff. The ability to target sanctions in multilateral games also implies that such programs have strategy options that are unavailable (or unused) by the programs described above. In effect, such programs may play "Tit-for-Tat" against *individual* team members, whereas the ordinary "Tit-for-Tat" player who responds to defection in kind can only engage in reciprocal behavior against the team as a whole. "Targeted Retribution" cooperates whenever it participates in a particular round of the game. It permanently exits from groups that include two defectors. In groups that contained a single defector in the previous production session, Targeted Retribution imposes a penalty of 2 units on the defector at a cost of 0.2 units to itself. If, in the following round, the entire group now cooperates, Targeted Retribution continues as a cooperator, otherwise it exits. Targeted Retribution is, thus, similar to the Moral Suasion program except for its recourse to a *strong and targeted* sanction.

3. Simulations: individual dispositions and joint production

Our round-robin tournaments forms all possible three "person" teams from the available behavioral programs. Scores are accumulated for each behavioral program and average payoffs per production session are computed. The relative performance of the behavioral programs in a particular tournament can be assessed by comparing their average payoffs per production session. The round-robin nature of the tournament implies that groups or teams are formed via an exogenous process. However, the teams that *continue* in the simulated scenarios developed below are self-selected in the sense that the *duration* of the team is determined by team member decisions to exit or not. Each team depends for its continued existence on their members' choices to participate in the joint enterprise. In this sense, continuing production teams may be considered to be *self-selected* even within our round-robin tournaments.

3.1. The value of exit in multilateral PD games: a tournament between five elementary programs

The five behavioral programs included in the first tournaments are a balanced group of cooperative and non-cooperative programs with two non-adaptive programs (Free-Rider, and Cooperator), two programs that exit (Prudence and Hit-and-Run), and Tit-for-Tat. None of these programs is "teachable" in the sense that a sanction will cause a non-cooperative strategy to become a cooperator. Thus, the success of cooperative strategies in this tournament is entirely based on their ability to end unproductive teams by leaving them and, thereby, benefiting from self-selection. The first series of tournaments explores the effects of changes in the potential duration of the joint enterprise on relative performance of these relatively simple strategies.

Fig. 1 summarizes the performance of the five behavioral programs in 20 round-robin tournaments. The maximum possible number of production sessions is listed on the horizontal axis, and average scores for each program are measured along the vertical axis. The potential duration of a joint enterprise is varied between one and twenty production sessions. The relative performance of cooperative and free-riding oriented programs clearly varies with the maximal number of production sessions (Fig. 1).

Overall, the average scores of all cooperation-oriented programs (Cooperator, Tit-for-Tat, Prudence) increase as the number of possible production sessions on a given team increases, while those of defection-oriented programs (Free-Rider, Hit-and-Run) decline as the number of possible production sessions increases.

Obviously, it is only in repeated games that the various adaptive strategies distinguish themselves. In very short-lived joint enterprises, the adaptive programs are essentially the same as the non-adaptive programs. Free-riding is a profitable strategy within such short lived enterprises. In settings where teams engage in fewer than three joint-production sessions, the strategies that begin with defection obtain the highest scores. Although the teams

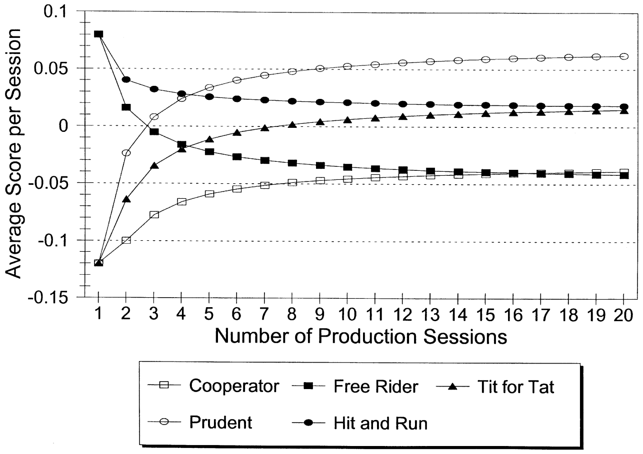

Fig. 1.

composed exclusively of defectors generate losses for all, those losses are more than offset by the gains that free-riders realize in groups with two cooperators. Conversely, teams composed exclusively of cooperative programs realize a mutual advantage from joint production in short-lived joint enterprises. Yet, those gains do not fully offset their losses from groups that include free-riders.

Adaptive programs generally do better than non-adaptive programs in PDE games when there is sufficient time for adaptation to take place. For example, even in a two production session per team setting, the adaptive strategies generally score better than the non-adaptive strategies. All the adaptive programs implicitly "punish" first round defection by others on the team. The two exiting strategies included in the first tournament exit from the groups in which some other players defect in the first round. This allows Hit-and-Run and Prudence to avoid continued losses from groups that include defectors. Tit-for-Tat continues to participate in such groups, but defects in period 2, i.e. it becomes a fellow free-rider. This response increases Tit-for-Tat's payoff and reduces the payoffs of *all other players* in continuing teams, including cooperators, which improves Tit-for-Tat's *relative* performance.

The importance of voluntary association is demonstrated by the relatively high scores of the two exiting programs in potentially long-term joint enterprises. Exit provides an effective and low-cost method of avoiding losses from poorly performing teams. Hit-and-Run and Prudence both outscore the other programs in tournaments with three or more joint production sessions. As the number of production sessions increases, the average scores of the conditional cooperative strategies increase as losses from the first and second rounds are increasingly offset by the gains from continued cooperation in profitable joint enterprises. The Prudence program, which never free-rides, is the most successful program in tournaments where joint production can continue for more than five production sessions.

Note also that it is conditional cooperation rather than cooperation per se that yields the best scores from participating in potentially long standing teams. Prudence is profitable when teams engage in at least three production sessions. The Tit-for-Tat program achieves positive net payoffs in tournaments with at least eight production sessions. The naive Cooperator program does not profit from even 15 session tournaments. This suggests that *the conditionality of cooperation and the possibility of exit are both important determinants of the viability* of cooperative programs. Both conditionally cooperative programs punish programs oriented toward defection in the sense that they reduce the performance of programs oriented toward defection relative to what they would have achieved with a program of unconditional cooperation.

Of course, the mere possibility of exit does not, itself, directly encourage cooperation in the short run. Exit also creates opportunities for opportunistic programs. Hit-and-Run dominates all other programs, except for Prudence, over the entire range of interest. Exit allows Hit-and-Run to score better than the unconditionally defecting program because exit enables Hit-and-Run to avoid other free-riders and the untargeted retribution of the Tit-for-Tat player. Hit-and-Run is dependent for its success in multiple encounter settings on the continuing exploitability of the unconditional Cooperator strategy, and the very limited information that other teams are assumed to initially have about its behavior.

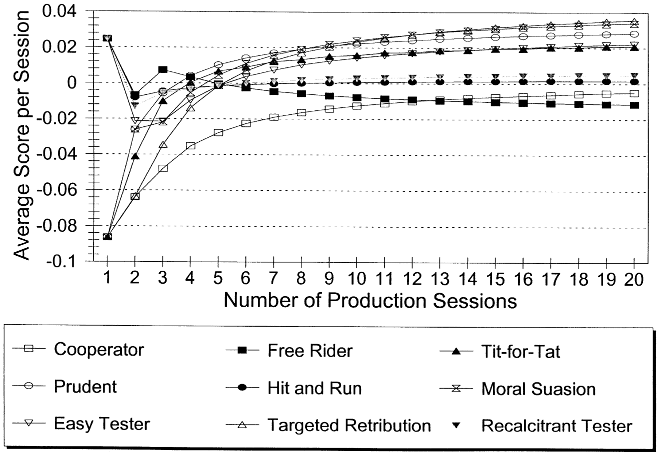

Fig. 2.

3.2. *Tournaments with teaching and teachable programs: the power of sanctions*

In the first tournament, sanctions could reduce the relative performance of programs oriented toward free-riding, but could not generate additional cooperation because there were no programs in the tournament that switch from free-riding to cooperation in response to sanctions. We now add two “teachable” programs: Easy-Tester and Recalcitrant-Tester, and two “teaching” programs to the joint enterprise tournament, Moral Suasion and Targeted Retribution. Easy-Tester is somewhat more willing to change from a strategy of defection to a program of cooperation than is Recalcitrant-Tester. Moral Suasion’s program uses exit as a method of changing the behavior of other team members. Targeted Retribution imposes, at a small cost to itself, a punishment on a free-rider in the previous round of joint production. (The Targeted Retribution program can be interpreted as an informal team manager with some enforcement authority or as an independent curmudgeon).[12]

Fig. 2 plots the results of a series of 20 tournaments including these nine programs. Notice that the best strategy for interacting with the teams in this population again varies with the potential longevity of each team. The payoff *envelop* is quite instructive. In joint enterprises with very short lives, the strategy of always defecting works very well. In somewhat longer lived teams, the fairly timid Prudent strategy works best, followed by Moral Suasion and, finally, Targeted Retribution. The latter two both bear individual costs for attempting to switch the behavior of other players. Targeted Retribution bears a higher cost in the short run but secures the larger benefit in the long run insofar as it is able to elicit cooperation from more types of teams in particular those including Recalcitrant-Tester and consequently is a member of more productive enterprises than is Moral Suasion (Fig. 2).

---

[12] In order to generate cooperation in multilateral PDs Boyd and Richerson (1992) relied upon a strategy of retribution or “*moralistic* strategies which cooperate, punish non-cooperators, and punish those who do not punish non-cooperators.” Our previous results indicate that such strategies are not required to achieve cooperative outcomes in multilateral PD games where exit is possible.

The ability to target sanctions at free-riders is important for teams, and for players in the long run, but again exit remains critically important. Exit provides a means of avoiding harm from non-cooperative players and also for punishing non-cooperators by ending joint enterprises which include such players. (The ability to target santions is assumed to be more costly in the short run than exit from voluntary enterprises). On intermediately lived teams, Targeted Retribution scores well, but not as well as the other cooperative exiting strategies. All four conditionally cooperative programs realize substantially greater payoffs than are realized by the strategy of unconditional cooperation. Moreover, only the "nicest" of the exiting free-riding programs, the Easy-Tester program, scores as well as the least effective conditionally cooperative program.

3.3. *Toleration and cooperation*

The difficulty of joint production in the previous settings made toleration a relatively costly activity. We now examine a setting where joint production is less demanding. We reduce the shared overhead cost from 2 units to 0.5 units. To keep the results of this tournament comparable to the previous ones, we retain the assumption that the exit of a single player yields a zero payoff for all team members, in spite of the fact that under the revised payoff structure two-person teams might now yield positive payoffs. Table 3 summarizes the payoffs that now accrue to cooperators and defectors in a single round of the joint enterprise game Table 3.

In settings where joint production is easier to undertake, tolerance may plausibly be greater insofar as private payoffs remain the most natural method of assessing team performance for an individual player. To analyze the effects of increased tolerance, we assume that the triggering thresholds for exit and Tit-for-Tat remain the same as in the first tournament. Note that this implies that the payoff-triggered programs are now more tolerant of free-riding than before, since a single free-rider will no longer generate a change in strategy. Targeted Retribution remains as intolerant as in the previous tournaments because it has a "number of defectors" based threshold. It still punishes a single defecting strategy, and exits from a joint enterprise if there are two defectors or if the targeted defector fails to subsequently cooperate.

Fig. 3 summarizes the average payoffs achieved in this setting. Note that the easier circumstances of joint production increase average payoffs for all the behavioral programs, while the greater tolerance of the conditional cooperators makes free-riding a much more successful strategy than in the previous tournaments. The two Tester programs now score

Table 3

Combination of strategies	Payoff for Cooperator	Payoff for Defector
All cooperate	0.833	–
Two cooperate, one defects	0.166	1.166
One cooperates, two defect	−0.5	0.5
All defect	–	−0.166
One or more Exit	0	0

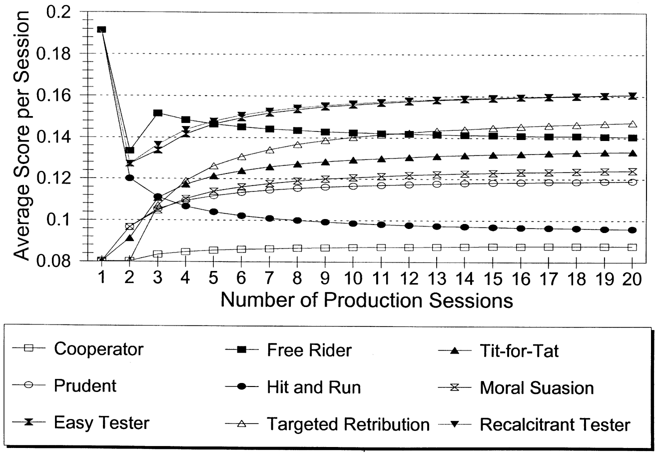

Fig. 3.

highest over most of the range of interest, followed by the Targeted Retribution program and the Free-Rider program.

Exit and conditional cooperation remain important determinants of average scores. The Tester strategies outscore the other strategies over much of the range of the tournament largely because they free-ride on most tolerant teams, but cooperate with the intolerant Targeted Retribution program, which exits from teams containing other more recalcitrant free-riders. The teachability of the Tester programs also improves the Targeted Retribution program's score. By opposing free-riding "in principle" Targeted Retribution obtains sufficient cooperation from Easy-Tester and Recalcitrant-Tester to more than compensate it for losses caused by its departures from other teams including uncooperative strategies (after punishing those Free-Rider programs).

It bears noting that toleration is crucial to the relatively good performance of the free-riding behavioral rules. Had we adjusted the conditional strategy thresholds from 0 to 0.5 (the payoff of solitary production under the revised payoff function), the relative performance of the programs would have been largely the same as in the previous tournaments. Additional exit and retaliation by conditional cooperators would have reduced the scores of free-riders and raised the average output of the teams that continue producing. Proclivities to cooperate and sanction both contribute to the relative success of the programs of conditional cooperation.

The assumed conditional strategy thresholds demonstrate that free-riding is most viable in settings where the fruits of cooperation are most easily obtained, insofar as it is in those circumstances that free-riders are most likely to be tolerated by conditional cooperators. Conversely, the response thresholds imply that the *greatest private advantage* from programs of conditional cooperation arise in ongoing joint enterprises where successful joint production requires significant cooperation. To the extent that the evolution of firms and other joint enterprises reflects efforts to realize more subtle and difficult to realize gains from joint production, the last two tournaments suggest that one would observe an increasing emphasis on cooperation in the work place. Larger and more specialized organizations would require increasingly cooperative personnel *organization men and women* possibly in conjunction with more elaborate incentive compatible contracts, to successfully realize the fruits of joint production.

4. Tournaments with endogenous populations

In the long term, variation in the relative performance of alternative behavioral programs for participating in joint enterprises may be expected to affect the relative frequency of those programs within a given population. The success and failures of others team members provide evidence that rational individuals can use to refine their own decision rules. That is to say, optimizing individuals are inclined to imitate the behavioral routines of their most successful peers. Moreover, even without conscious efforts to imitate successful behavioral programs, evolutionary pressures also tend to promote behavioral routines that have larger payoffs and thereby a social and biological survival advantage in the long run.[13] Both processes lead one to expect the relative frequency of the higher scoring programs to increase and those of the lower scoring programs to decrease through time as imitation and evolution take place.[14] On the other hand, evolutionary processes are not always as straightforward as intuition suggests because the relative performance of the behavioral programs is affected by changes in the *composition* of the population from which joint enterprises are formed.

To gain some insight into the population dynamics of a society facing joint production problems, we modify the simulation program to allow the relative frequency of our nine behavioral programs to vary systematically through time. We use a truncated version of what has come to be called the replicator dynamic. We assume that the total population of each “generation” of potential team members is constant, but that the relative frequency or population shares of the available behavioral programs varies with their past performance. A constant population size increases the comparability of successive generations and implicitly increases the intensity of competition between programs insofar as it implies the existence of relevant economic scarcities. The composition of each generation reflects the relative performance of the behavioral programs in the previous round-robin tournament.

The simulation software computes base scores for each member of each team, B_{ijk}, and weights those base scores by the relative frequency of specific teams in generation t to calculate each program’s expected payoff. The relative frequency of a particular *team*, *ijk*, in a particular generation, t, is the product of the relative frequencies of programs on that

---

[13] Simulations based on genetic algorithms are widely used to study evolutionary processes. We do not rely upon the GA methodology here, because of the difficulty of interpreting and presenting the results generated. For example, in a five session PDE game there are about a thousand strategic sequences that must be represented and interpreted ($2^5 \times 2^5$: five cooperate or defect decisions combined with five play or not decisions). Even with the relatively small number of behavioral programs examined and small production teams simulated, there are generally more than a hundred teams (joint enterprises) in each of the simulated tournaments examined. The large number of teams and incentive structures examined effectively rules out a single uniform experimental and analytical examination of the relative performance of these programs in the wide variety of settings explored.

[14] Witt, 1986, and others, regards this to be a process of social learning. Although all learning takes place within an individual’s mind, each individual learns from the successes and failures of other individuals that he observes or communicates with. This extends each individual’s “knowledge” beyond his own direct experience and invention.

team, $f_{it}f_{jt}f_{kt}$. The expected payoff or "cumulative score" of program i in generation t, S_{it}, is thus a weighted average of its base payoff from all teams on which it is a member.

$$S_{it} = \sum_{j}\sum_{k} B_{ijk}(f_{it}f_{jt}f_{kt}). \tag{3}$$

If program i that accumulates twice the score of program j in generation $t-1$, it will have a relative frequency twice as large as that program in generation t, $f_{it} = 2f_{jt}$, as long as both scores are greater than zero. Programs that accumulate negative scores, $S_{it} < 0$, are deemed non-viable, and do not survive into the next generation. The absence of such programs may be interpreted as a general lack of interest in imitating unprofitable behavioral programs, or as a consequence of the fundamental biological or economic non-viability of routines that yield negative net returns. However, because the performance of a behavioral program is population dependent, the simulations allow eliminated programs to "reemerge" in future generations if the population mix becomes favorable to their particular approach to joint production games. The possibility of reemergence, as we will see below, allows defection prone strategies to almost continually threaten populations of cooperators.

To determine whether a program has returned to viability, "shadow scores" are calculated for each eliminated program using a population weight of 0.01 and the actual population weights of all the *surviving* programs. (Absent programs may, thus, implicitly form teams with themselves, but not with other unused programs). If an eliminated behavioral program would have achieved a positive score, the program *reenters* the population in the next generation with a population share based on its shadow score. Eliminated behavioral strategies are thus assumed to be more like "crab grass" than "dinosaurs" even very bad behavioral routines rarely vanish forever.

4.1. *Population dynamics in five session teams: evolutionary stability*

Fig. 4 characterizes the variation in the populations shares of the feasible behavioral programs over twenty generations in a relatively tough environment for joint enterprises. Each generation participates in a round-robin joint production game based on one reported in Section 3. We assume that each team can engage in joint production at most *five* times, and that each program initially is equally represented in the population. Five production sessions is the shortest that allows all the various sequences of adaptation to be played out. Average scores in the first generation are, consequently, those reported in Fig. 2 for five sessions.

Note that only Prudent, Moral Suasion, and Tit-for-Tat secure positive scores in the first generation. So, the second generation is one of extreme adjustment. The poor performance of the defection-oriented programs in a five-session tournament largely reflects their interaction with *each other*, and the various sanctioning responses of the conditionally cooperative strategies. The second generation includes *only* cooperative strategies, and is therefore much more hospitable to defecting programs, albeit largely in the early sessions of each team. The possibility of "reentry" implies that completely cooperative societies can not be sustained. All of the eliminated strategies return to the pool of potential team members in the third generation.

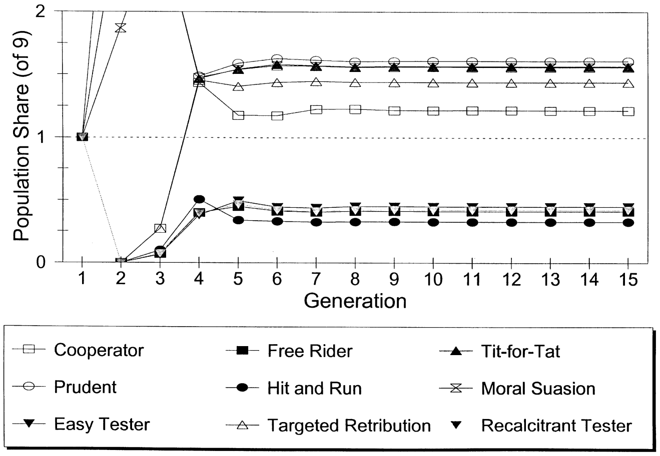

Fig. 4.

The society rapidly achieves stable population shares. After the fifth generation, population shares remain steady, that is to say, a fixed point is reached. The relatively cooperative environment that emerges allows a limited number of defection-oriented programs to survive, although even the unconditionally cooperative strategy achieves a larger population share than the four defection strategies. As often the case in a Hawk–Dove equilibrium, the steady state period exhibits a clear pattern with *all* the cooperative strategies having above average population shares, and all the uncooperative programs having below average population shares. Even in the very long run, one may observe both defection and cooperation in the joint production environment simulated.[15]

In the present environment, the cooperative strategies outscore the free-riding strategies largely because the main *sustained* source of positive benefits in the game arises from cooperation. The free-riding programs remain viable, although less profitable as a group, because of their success in the early production sessions, and because of the continuing *unconditional* cooperation of the Cooperator program.

---

[15] Stable population shares requires $F_{it} = F_{it+1}$, for all i strategies present in the population. This occurs when $S_{it} = S_{it+1}$ for each strategy. Under our assumed dynamic, the latter assures that $F_{it}/F_{kt} = F_{it} + 1/F_{kt} + 1$ since $F_{it+1} = S_{it}/\sum S_{jt}$. For example, stable populations shares occur when if scores stabalize at $S_{it} = \alpha S_{kt}$ so that $F_{it+1} = \alpha F_{kt+1}$. Clearly, a can be any real number greater than zero. On the other hand, recall that $S_{it} = F_{it}\sum_k\sum_j F_{kt}F_{jt}B_{ijk}$. Thus, $F_{it+1} = S_{it}/\sum S_{jt}$ can be rewritten as $F_{it+1} = F_{it}\sum_k\sum_j F_{kt}F_{jt}B_{ijk}/\sum_i\sum_k\sum_j F_{kt}F_{jt}B_{ijk}$, which implies that $F_{it+1}/F_{it} = \sum_k\sum_j F_{kt}F_{jt}B_{ijk}/\sum_i\sum_k\sum_j F_{kt}F_{jt}B_{ijk}$. The denominator of the right-hand side can be interpreted as the average payoff to a representative strategy, and the numerator that of a particular strategy. If $F_{it+1} = F_{it}$, each strategy must earn the average payoff, e.g. realize the same profit, in the case where stable population shares are reached. Thanks are owed to one referee and to Robin Hanson for making this point clear to us. In the present context, stability emerges when there are enough cooperators to allow non-cooperative strategies to realize the same profit from a large number of short lived teams as cooperators do on longer term teams net of costs imposed on them by defectors. Of course, in some parameterizations of the production environment, one or both classes of strategies may completely disappear insofar as average payoffs are never equalized, as in some cases below.

As in the static simulations, Prudence is the most successful program in five session evolutionary tournaments. The relatively short life span of the teams and the relatively low frequency of the teachable Easy-Tester programs causes Moral Suasion and Targeted Retribution to be somewhat less successful than the more timid Prudent program. Hit-and-Run initially scores lowest and is the least frequent program, while the Tester programs are slightly more common than their unteachable counterparts. However, neither Tester program did as well in the three-person setting as in the two-person evolutionary simulations of our previous paper where Easy-Tester was second only to the Moral Suasion program.

The particular path followed and population shares obtained vary systematically with minor changes in program parameters. For example, as the number of sessions increases, the population shares of the free-riding oriented programs all fall. The basic result that cooperative programs outperform free-rider oriented programs is robust in settings where production teams are reasonably long lived and exit is possible. Although significant differences occur, the results are broadly similar to the two-person PDE evolutionary results reported in Vanberg and Congleton (1992). The new results suggest that targeted signals/sanctions are not critically important to the emergence of cooperation within small joint enterprises in an environment where exit is possible.

4.2. Population dynamics on three session teams: cyclic anarchy?

The static simulation results reported in Fig. 3 indicate that the cooperative programs are less successful in settings where few production sessions take place. This suggests that the evolutionary path of such games will be less advantageous for programs of conditional cooperation and possibly less well behaved than in the five encounter setting examined above. We now examine the evolutionary path of the same nine teams in an even tougher setting for joint production, one where production teams may continue for at most three sessions.

Fig. 5 depicts the ensuing dysfunctional and erratic evolutionary path. As indicated by Fig. 3, the first generation performance of a three session per team setting is such that only the

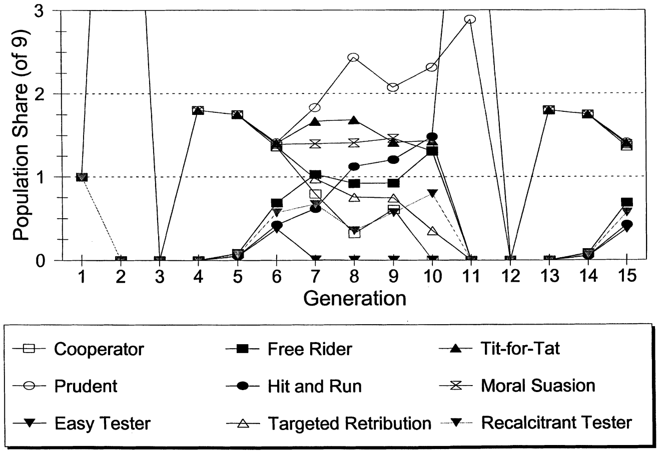

Fig. 5.

free-rider program obtains a positive payoff. Consequently, the second generation consists entirely of players using the Free-Rider program. Joint production is not viable in such a community. Consequently, *none of the programs* survive into the third generation. In this tabula rasa setting, the simulation program computes shadow scores for teams with identical team mates, e. g. family enterprises. Consequently, the cooperative programs all re-emerge. The preponderance of cooperative strategies in the fourth generation, in turn, makes all the other strategies potentially viable leading to their reemergence in generation 5. The following generations evince a gradual increase in the number of free-riding programs and a decline in the cooperative programs which favors the timid "Hit-and-Run" and "Prudent" programs. However, in generation eleven, the preponderance of the Hit-and-Run program causes negative payoffs for both surviving programs and for all of the unused programs. Joint production disappears in generation twelve, which allows the reemergence of the cooperative programs in generation thirteen. This begins another production and population cycle.[16]

4.3. Population dynamics with a powerful Targeted Retribution player

The social dilemma faced in the previous evolutionary simulation is that the success of cooperative programs caused an overwhelming reemergence of the defection oriented programs. When there are only three encounters in each group, the conditionally cooperative strategies do not sufficiently benefit from the formation of ongoing joint enterprises with fellow cooperators to offset losses from dealings with non-cooperators, nor do they sufficiently penalize the free-riding strategies to reduce their numbers in future generations. To determine whether the former or the latter is the main factor leading to the non-viability of the cooperative programs in this setting, we now modify the Targeted Retribution program so that it is able to impose a stronger sanction on defectors. We assume that Targeted Retribution imposes a cost of 3 rather than 2 on single defectors. In effect, teams that include Targeted Retribution now have a strong team manager empowered to severely punish a free-rider. This penalty significantly reduces the net benefits that accrue to free-riding programs on teams that include a Targeted Retribution player. Fig. 6 depicts the results of a three production session evolutionary simulation with the described modification of the Targeted Retribution program. Note that the more stringent penalty imposed by Targeted Retribution in combination with the untargeted sanctions of the other players reduces the payoffs from free-riding sufficiently to stabilize the community. In fact, the new results look very much like those of the five production session population path of Fig. 4. Cooperative strategies all have above average population shares and free-riding strategies all have below average shares. Although this more benign result is entirely attributable to Targeted Retribution's effort at the margin, the latter is not the most common strategy in the new population Fig. 6.

---

[16] The small probability that each non-viable strategy finds itself on a team with other players like itself allows joint enterprise to re-emerge in the 4th, 13th, and 22nd generations, but not to flourish in the long run. Initially, small clusters of cooperative programs reemerge, but the eventual reemergence of the defection inclined programs generates a population mix where all programs again yield negative scores after a few generations. In our simulations, this pattern would repeat itself ad infinitum.

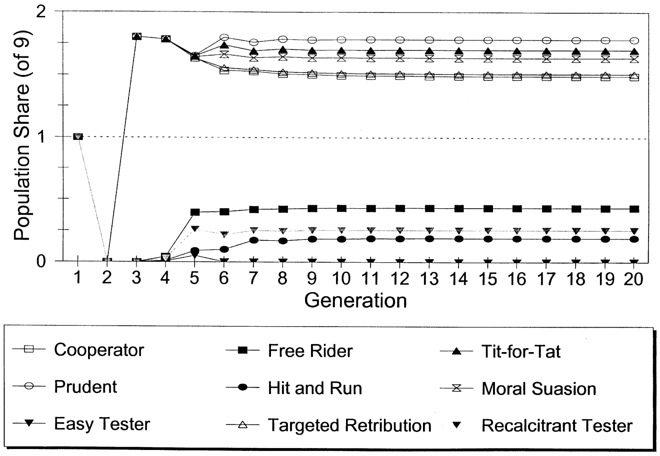

Fig. 6.

The relatively better performance of the other conditionally cooperative programs occurs for two reasons. First, all of the cooperative programs benefit as the number of defection oriented programs diminishes. Targeted Retribution bears the cost of imposing effective sanctions, but does not realize all of the benefits generated. Second, as the number of Easy-Tester and Recalcitrant-Tester programs falls, the relative performance of Targeted Retribution diminishes since it has fewer teachable team mates who can be transformed into cooperative team members.

Although Targeted Retribution is not the most profitable strategy, it is sufficiently profitable to survive in the long run and to stabilize joint production in the settings explored here. In settings where the costs of Targeted Retribution are higher than assumed here, some new form of compensation may be necessary to secure this “policing” or “managerial” behavior. For example, joint enterprises may have to hire a constable or joint enterprise manager in order to realize the advantages of joint production.[17]

5. Conclusion

Our analysis has focused on the relative performance of nine broadly representative behavioral routines, or heuristics, for participating in small voluntary joint enterprises. The exploitative and of cooperative strategies that we have analyzed represent a fundamental set of behavioral programs that can be refined or combined in numerous ways to make more complex strategies. Yet, even the relatively modest enterprise of examining the performance

---

[17] R. Axelrod’s The Evolution of Cooperation (1984) and his later contribution on “An Evolutionary Approach to Norms (1986) also explore the respective roles of strategies of reciprocity and of Targeted Retribution discussed here. In particular the 1986 study is concerned with multilateral PDs and a critical role is played by assumptions about “vengefulness” as a disposition to punish cheaters as well as those who fail to punish cheaters. See also (Coleman, 1990).

of these basic programs for participating in multilateral PDE settings generates instructive insights.

Our particular interest has been in the performance of non-defecting programs, i.e. programs that cooperate whenever they play, but avoid defectors, and in the importance of sanctions for the success of joint enterprises. We find that the possibility of exit allows individuals to escape from dysfunctional teams which increases the potential benefits of cooperation while reducing those associated with free-riding. In long-term enterprises, the simulations indicate that the sorting that exit allows is sufficient to generate a large number of viable joint enterprises composed of conditional cooperators. Punishment, especially targeted punishment, is more critical to the success of short-lived enterprises. In short-term enterprises, and in enterprises where cooperating players are more tolerant of free-riders, programs oriented toward free-riding were more successful than cooperative strategies, and consequently fewer resources were invested in joint production overall.

Similar results held in evolutionary settings where successful behavioral programs may be imitated, or propagated by non-rational biological or social processes. In shorter lived joint enterprises, cooperative strategies were rapidly replaced by strategies inclined toward defection and joint enterprise disappeared. Such settings exhibited the dysfunctional pattern of joint production implied by elementary game theory and public goods theory. However, the existence of a single cooperative program with the power to penalize free-riding can be sufficient to restore the competitive edge to programs of conditional cooperation and thereby to facilitate joint production. To the extent that civil societies have such players, or have evolved institutional equivalents, programs of conditional cooperation may remain personally rewarding rational strategies even in short-lived joint enterprises.

In evaluating the successful performance of the cooperation-inclined programs in our simulations, one should take into account that the setup of the tournament tended to favor defection-inclined programs because of the nature of the joint-production process and the absence of reputation effects and/or any ability to initially recognize the behavioral types of fellow team members. Any positive degree of type-recognition would work in favor of cooperators, and in a world in which reputation plays a role in social interaction we can plausibly assume that players would not choose their partners for joint enterprises entirely blindly, Frank (1987, 1988). However, our results demonstrate that the ability to assess *team* performance, together with the possibility of exiting from undesirable teams, can be sufficient to allow conditional cooperators to outperform their defection inclined counterparts in all but the most short-lived joint-production enterprises. Altruism is not a necessary precondition for cooperation in PDE settings. However, non-cooperative programs do not disappear even in the long run, but rather linger on at the margins of the more civil society as somewhat less attractive alternatives to strategies of conditional cooperation.

In an economic context, our results suggest that the cooperative tendencies of a significant subset of potential employees can potentially be relied upon as a method for solving free-riding and other prisoner's-dilemma problems within firms and other organizations. That is to say, hiring employees with the "right values" can potentially enhance an organization's efficiency. There are, of course, other solutions to the PD-type problems within firms and other voluntary organizations. For example, team-members can adopt contractual devices that appropriately change the incentives faced by individual participants. Or, equivalently, each game of joint production could be embedded in a super-game where reputation or status losses encourage cooperation by "local" net-return maximizers.[18] On the other hand, when cooperative types of potential employees exist, selecting the "right" production teams becomes an alternative method of addressing free-riding problems.

The importance of having the "right people" in the "right places" is clearly evidenced in practice by the extensive screening efforts carried out by both profit and non-profit organizations for essentially all senior and junior employees. Our results suggest that firms and other voluntary organizations more widely rely upon more or less cooperative individuals to solve problems of joint production in the long run than is generally acknowledged by economic theory. Completely incentive compatible contracts are clearly less critical in a setting where it is possible to hire "team players" (conditional cooperators) and "effective managers" (targeted retributors).

The viability of the strategies that exclude defection from their behavioral repertoire is also relevant for research beyond economics. If such strategies were not viable, rationality, in the sense used here, could not account for the many occasions where such behavior was observed, and one would not expect to see them routinely exercised in real or experimental multilateral PDE settings. Our results suggest that non-defecting programs of conditional cooperation can be rational average payoff maximizing strategies for participating in small multilateral PDEs in a wide variety of settings, partly because groups whose members employ such strategies are able to mutually profit from cooperation, and partly because players who use such strategies punish free-riders.

Low cost exit provides conditionally cooperative players with a means for avoiding exploitation in non-cooperative groups, for punishing non-cooperators, and for retaining relations with well-functioning groups. In this manner, societies with numerous opportunities for voluntary association as typical of most modern market based societies facilitate the formation of cooperative clusters and the sorting of cooperative and non-cooperative players. The conditional propensity to cooperate and punish can be *rational* dispositions that serve the person well who adopts them in such settings — as well as the joint enterprises that include and retain such persons.

References

Alchian, A., 1987. Property Rights. In: The New Palgrave, Vol. 3. pp. 1031–1034.

Alchian, A., Demsetz, H., 1972. Production, information costs, and economic organization. American Economic Review 62, 777–795.

Andreoni, J., Miller, J., 1993. Rational cooperation in the finitely repeated prisoner's dilemma: experimental evidence. Economic Journal 103, 570–585.

Axelrod, R., 1984. The Evolution of Cooperation. Basic Books, New York.

Axelrod, R., 1986. An evolutionary approach to norms. American Political Science Review 80, 1095–1111.

Boyd, R., Richerson, P.J., 1992. Punishment allows the evolution of cooperation (or anything else) in sizable groups. Ethology and Sociobiology 13, 171–195.

Buchanan, J.M., 1965. An economic theory of clubs. Econometrica 32, 1–14.

Coleman, J.S., 1990. Norm-generating structures. In: Cook, K.S., Levi, M. (Eds.), The Limits of Rationality. The University of Chicago Press, Chicago and London, pp. 250–273.

---

[18] For a nice overview of incentive compatible contractual mechanisms see Mas-Colell et al. (1995, Chapter 14). See Congleton (1989) for a discussion of efficiency enhancing status games. Congleton (1991) discusses the efficiency enhancing role that a work ethic can play in a market economy.

Congleton, R.D., 1989. Efficient status seeking: externalities, and the evolution of status games. Journal of Economic Behavior and Organization 11, 175–190.

Congleton, R.D., 1991. The economic role of a work ethic. Journal of Economic Behavior and Organization 15, 365–385.

Cooper, R., Dejong, D., Forsythe, R., Ross, T., 1996. Cooperation without reputation: experimental evidence from prisoner's-dilemma games. Games and Economic Behavior 12, 187–218.

Frank, R.H., 1987. If Homo Economicus could choose his own utility function, would he want one with a conscience? American Economic Review 77, 593–604.

Frank, R.H., 1988. Passions Within Reason: The Strategic Role of Emotions. W.W. Norton, New York and London.

Hayek, F.A., 1978. Nature v. Nurture Once Again, In: New Studies. University of Chicago Press, Chicago (Chapter 19).

Hirshleifer, D., Rasmusen, E., 1989. Cooperation in a repeated prisoners'-dilemma with ostracism. Journal of Economic Behavior and Organization 12, 87–106.

Hirshleifer, J., Coll, J.C.M., 1991. The limits of reciprocity, rationality. Rationality and Society 3, 35–64.

Kreps, D.M., 1990. A Course in Microeconomic Theory. Princeton University Press. Princeton, NJ.

Kreps, D.M., Migrom, P., Roberts, J., Wilson, R., 1982. Rational cooperation in the finitely repeated prisoner-dilemma. Journal of Economic Theory 27, 245–252.

Mailath, G., 1998. Do people play nash equilibrium? Lessons from evolutionary game theory. Journal of Economic Literature 36, 1347–1374.

Mas-Colell, A., Whinston, M.D., Green, J.R., 1995. Microeconomic Theory. Oxford University Press, New York.

Nalbantian, H.R., Schotter, A., 1997. Productivity under group incentives: an experimental study. American Economic Review 87, 314–341.

Rubin, P., 1982. Evolved ethics, evolved ethics and efficient ethics. Journal of Economic Behavior and Organization 3, 161–174.

Simon, H.A., 1979. Rational decision making in business organizations. American Economic Review 69, 493–513.

Sugden, R., 1986. The Economics of Rights, Cooperation and Welfare. Blackwell, Oxford.

Vanberg, V., Congleton, R.D., 1992. Rationality, morality and exit. American Political Science Review 86, 418–431.

Witt, U., 1986. Evolution, evolution and stability of cooperation without enforceable contracts. Kyklos 39, 245–266.