

To abide MIT's open-access policy and also by copyright issues, this is a pre-publication version that contains the same content as the journal publication. We apologize for any typographic mistakes. The journal citation is: Fader, Peter and John R. Hauser (1988), "Implicit Coalitions in a Generalized Prisoner's Dilemma," Journal of Conflict Resolution, 32, 3, (September), 553-582. <http://doi.org/10.1177/0022002788032003008>

Implicit Coalitions in a Generalized Prisoner's Dilemma

PETER S. FADER

*The Wharton School
University of Pennsylvania*

JOHN R. HAUSER

*Sloan School of Management
Massachusetts Institute of Technology*

The presence of a third party can affect attempts by two players to cooperate in a three-player, continuous-alternative, repeated Prisoner's Dilemma-like game. If the third player is uncooperative, two players may find it advantageous to cooperate implicitly, at a level somewhere between full (i.e., three-way) cooperation and full defection. We examine this phenomenon of implicit coalitions via two sequential computer tournaments (38 algorithms in tourney 1, 44 algorithms in tourney 2). In both tournaments, each with a different payoff function, the ability to recognize and / or encourage implicit coalitions seems to be a key indicator of success. This result holds up in a test of robustness. We also examine other properties, including those identified earlier by Axelrod (1980a, 1980b).

MORE THAN TWO PLAYERS

At a time when the superpowers appear to be moving toward some degree of cooperation on nuclear weapons, there is a growing concern about the nuclear capability among a number of nations in the Middle East (Barnaby, 1987). One concern is whether the presence of a noncooperating outside player will encourage or discourage cooperation among the superpowers. It is not yet clear how these outside players will affect the level of cooperation the superpowers might achieve.

Consider the dramatic effect of outside players on OPEC. After a decade of highly profitable cooperation (collusion) the cartel collapsed, partially because of increased production by non-OPEC nations such as the United Kingdom. Member nations began to cheat more and more from the agreed price and production guidelines, reaching a climax in 1986 when Saudi Arabia was no longer willing to be the sole co-operator. Early in 1986, the Saudis finally gave up their attempts to maintain cooperation and began increasing output and offering price discounts in an attempt to "punish" the United Kingdom. Oil prices dropped as low as \$8/ barrel. Recently, the Saudis and OPEC have attempted to stabilize prices at a more moderate level than before, finally acknowledging the critical role of outside nations.

In another example, firms in the U.S. microelectronics industries, adversely affected by the growing influence and economic power of foreign competition, have formed the Microelectronics and Computer Technology Corporation to cooperate on basic and applied research (Griffin, 1987). Although members risk loss of competitive research advantage relative to other U.S. firms, the potential gains may well justify the almost \$50 million investment.

Whether or not it is in the social good to encourage such coalitions, it is important to understand how coalitions can influence the development of effective strategies in games

involving more than two players. This article addresses the role implicit coalitions in a repeated, generalized Prisoner's Dilemma (GPD). The classical Prisoner's Dilemma, a two player, two-act game, captures the essential conflict between unilateral incentives (i.e., more sales through price discounts) and group incentives (price restraint and higher profits). The GPD extends this basic form of conflict to a richer, more complicated setting—an N -player game with many possible actions, either discrete (e.g., number of warheads) or continuous (e.g., price levels or R&D investment). After formalizing the GPD and motivating implicit coalitions, we describe two competitive strategy tournaments in the spirit of Axelrod (1980a, 1980b). Results of the tournaments illustrate the importance of implicit coalitions in a repeated GPD. We describe one strategy that seems to encourage coalitions and we test its robustness across a series of environments that vary from very "nice" to very "nasty."

FORMALIZATION

CLASSICAL PRISONER'S DILEMMA

Over its 30-year lifespan, the PD has been one of the most frequently studied phenomena in economics, political science, sociology, and psychology. (See Axelrod, 1984, for a review of these and other applications of the PD.) The classic 2 X 2 PD allows each player to either cooperate (C) or defect (D). If both players cooperate in a given period, then each is rewarded with a payoff of r points. If one player defects from mutual cooperation, she or he receives the temptation payoff of t , while the cooperating player gets the "sucker's payoff" of s . If both choose to defect, then each receives the punishment payoff of p .

		Player 2	
		C ₂	D ₂
Player 1	C ₁	$r_1=3$ $r_2=3$	$S_1=0$ $t_2=5$
	D ₁	$t_1=5$ $S_2=0$	$P_1=1$ $P_2=1$

Figure 1: Classic Prisoner's Dilemma. Subscripts on actions (C, D) and payoffs (t, r, p, s) indicate the player.

Figure 1 illustrates a typical set of payoffs for the 2 X 2 PD. A quick scan of Figure 1 reveals that each player has the unilateral incentive to defect, regardless of the other player's decision, but if the two players cooperate, both achieve high scores. Thus, if this were a one-shot game, each would be best off defecting since there would be no incentive to deviate from that action. However, if the game were repeated, strategies might change.

In general, the PD property holds if the payoffs r , t , s , and p must meet certain constraints. The essential property, once again, is that each player has a dominant alternative (to defect), but if both defect, the resulting payoff (p) is less than the payoff for mutual cooperation (r). Specifically,

- (1) Regardless of what our opponent does, we are best off defecting. If she or he cooperates we prefer to defect (i.e., $t > r$), and if she or he defects we still favor defection ($p > s$)

- (2) Regardless of what option we choose, we are better off if our opponent is lenient and chooses a dominated alternative (cooperation). Thus $r > s$ for then we cooperate, and $t > P$ For when we defect.
- (3) Mutual cooperation is always preferred to mutual defection: $r > p$.

These three sets of inequalities can be combined into one compound inequality: $t > r > p > s$. This is the heart of the PD. When the game is repeated, some researchers add a fourth condition to discourage oscillations:

- (4) Continued cooperation is better than alternating between cooperation and defection: $2r > t + s$.

GENERALIZED PRISONER'S DILEMMA (GPD)

In order to study implicit coalitions among N players we use the PD framework to balance unilateral incentives with group cooperation. We do not claim that all conflict situations are PD's; we claim only that many interesting situations are consistent with the PD paradigm. Thus we state a set of conditions that apply to N players and that, in the case of two players, reduce to the classical PD conditions.

Several researchers have proposed and analyzed N -player PDs, usually to study the behavior of large groups or entire communities. For example, each person finds it easier to litter than to carry paper to a wastebasket, but society as a whole is better off if no one litters. This scenario is often referred to as the "Tragedy of the Commons," first proposed by Hardin (1968). (See also Hamburger, 1979; Goehring and Kahan, 1976; Taylor, 1976; Dawes, 1980; and Schelling, 1973.) The primary mode of analysis for many players is the payoff functions $C(n)$ and $D(n)$, which describe the payoffs to each cooperator and each defector when exactly n parties cooperate.

While these N -player models are an excellent way to study situations involving many players facing a binary alternative, we are more concerned with games involving fewer players and more alternatives. For example, we wish to study how two cooperators should respond to a defector when all three players have a continuous range of alternatives available to them. We seek to determine, among other things, whether they should continue to cooperate, switch to defection, or take some action between cooperation and defection. To address these issues our generalization must deal with continuous actions.

We define the GPD in terms of payoffs, P , and actions, A . In particular, let $P_i(A_1, A_2, \dots, A_N)$ be the payoff to player i if the N players take actions A_1 through A_N . We assume the payoffs are symmetric.¹ We define two key actions, the short-term, noncooperative, payoff-maximizing action, A^d , and the joint-payoff-maximizing action, A^c . A^d and A^c correspond to D and C in the classical PD. Each player can maximize its payoffs by choosing A^d , regardless of the actions of the other players. At the other extreme, if all players cooperate and choose the same action, A^c will maximize joint payoffs.

For simplicity of exposition we first consider games in which A^d is fixed and invariant with respect to competitors' actions. We relax this assumption in the second tournament. In some games, cooperation means less action—fewer weapons, less quantity produced, or less aggression. In other games, cooperation means more action—more missiles removed from

¹ This is not a critical assumption. For example, positive linear transformations that vary by player do not affect our analysis. By symmetric we mean $P_i(\dots, A_i, \dots, A_i) = P_j(\dots, A_j, \dots, A_j, \dots)$.

the European theater, higher prices, or more joint research. Without loss of generality, we consider the latter class of games and assume A^c is less than A^d . Therefore, “high” action is taken to mean more cooperative action throughout this article.

We now define a GPD. Because our players are symmetric, we state the conditions for player 1. It is understood that each condition applies to all players.

- (1) As long as $A_1 > A^d$, player 1 increases its short-term payoff by defecting further:

$$\frac{\partial P_1}{\partial A_1} < 0 \quad \text{for } A_1 > A^d$$

An alternative interpretation is that unilateral movement toward cooperation decreases payoffs. This condition generalizes $t > r$ and $p > s$ in the classical PD. Note that by the definition of A^d as the payoff maximizing action, we implicitly assume that payoffs decrease when actions are decreased below A^d .

- (2) Any move toward unilateral cooperation by an opponent increases the payoff to player 1:

$$\frac{\partial P_1}{\partial A_j} > 0 \quad \text{for } j=2, 3 \dots N$$

This condition generalizes $r > s$ and $t > p$. Note that it applies for all feasible actions by all of player 1's competitors.

- (3) Mutual cooperation is profitable. If all players increase their actions by the same amount, all are better off (as long as no actions exceed A^c).

$$\frac{\partial P_1}{\partial A_1} + \frac{\partial P_1}{\partial A_2} + \dots + \frac{\partial P_1}{\partial A_N} > 0 \quad A_j < A^c \text{ for all } j$$

This condition generalizes $r > p$. Note that we have defined condition 3 for all actions, not just symmetric actions and, by the definition of A as the joint-payoff maximizing action, we have assumed implicitly that condition 3 reverses for all A_j above A^c .

- (4) We wish to rule out profitable oscillations in an analogy to $2r > i + s$. There are many possible generalization to this condition; we choose a simple one by making it unattractive to take turns reducing actions unilaterally. That is,

$$\begin{aligned} & \frac{\partial}{\partial \alpha} [P_1(A_1 - \alpha, A_2, \dots, A_N) + P_1(A_1, A_2 - \alpha, \dots, A_N) + \dots + P_1(A_1, A_2, \dots, A_N - \alpha)] \\ & < 0 \text{ for } A^d \leq A_j - \alpha, A_j \leq A^c, j = 1, 2, \dots, N. \end{aligned}$$

AN EXAMPLE

Our first tournament was framed in terms of a triopoly in which scores correspond to profits and the actions are prices. For realism we chose the commonly used “constant elasticity” model of consumer response to prices in a differentiated triopoly. The parameters of the model, the elasticities, were chosen to be consistent with empirical estimates for a

variety of markets (e.g., Telser, 1962; Lambin, Naert, and Bultex., 1975; Lambin, 1976; Simon, 1979). We assumed “constant returns to scale” and chose scaling constants so the payoffs were easy to understand.

Specifically, the payoff function we used was (in terms of player 1):

$$P_1 = 3375A_1^{-3.5}A_2^{2.5}A_3^{2.5}(A_1 - 1) - 480 \quad (1)$$

A little calculus (taking the derivative of equation 1 and setting it equal to zero) yields the noncooperative payoff-maximizing action, $A^d = 1.40$, which is independent of competitive actions. By assuming $A_1 = A_2 = A_3$, we can solve for the cooperative action, $A^c = 1.50$. The reader can verify that conditions 1 through 4 hold for equation 1.

		Player 2 and 3	
		$A_2=A_3=A^c$	$A_2=A_3=A^d$
Player 1	$A_1=A^c$	$P_1=20$ $P_2=P_3=20$	$P_1=3$ $P_2=P_3=21$
	$A_1=A^d$	$P_1=29$ $P_2=P_3=11$	$P_1=12$ $P_2=P_3=12$

Figure 2: Example Payoffs When players 2 and 3 Choose Identical Actions.

When we restrict actions to A^d and A^c , the game defined by equation 1 becomes the classical PD. Suppose, for the sake of illustration, that players 2 and 3 are committed to choose the same action as each other. Figure 2 shows the possible pay offs under this restriction. These payoffs, although asymmetric because $P_1(A^d, A^c, A^c) \neq P_1(A^d, A^d, A^c)$, clearly obey the constraints for the classical PD. Of course, this restriction on players 2 and 3 is not realistic, nor is it imposed in the tournaments. Figure 2 simply illustrates the close relationship between the classical PD and the GPD.

Before proceeding to our analysis of implicit coalitions, we note one more important feature of the GPD model, the *envious* price. Many researchers have noted that human players in experimental PD games often defect in an attempt to beat their rivals rather than to score well for themselves. In the GPD, a distinct action, A^e , is associated with this type of behavior. The envious action is defined as the action that maximizes one player’s share of total payoffs. It is consistent with the notion of *difference maximization* as discussed by Shubik (1959). Any player who misses the main point of the game (i.e., maximize own score) and instead plays to maximize share of total payoffs will frequently choose the envious action. In the game based on equation 1, the envious action is calculated to be $A^e = 15/11 \approx 1.36$. In an oligopoly, managers might choose an envious action if they are rewarded on outcomes relative to other firms in the industry (i.e., bonuses based on market share). Note that in the short run, players rarely have a legitimate incentive to choose A^e . They can always do better for themselves (in

a single period) by raising actions from A^c to A^d . However, A^c might prove useful as a severe punishment for noncooperative behavior.²

IMPLICIT COALITIONS

Suppose player 3 in a repeated three-player game has chosen a strategy of consistently choosing A^d . How should players 1 and 2 react?

One option is to punish the defector by reciprocating her or his totally *noncooperative behavior*. For example, Axelrod (1981) showed that this type of strategy (ALL- A^d) is a best response to itself in the two-player game. But in the multiplayer game, defection in response to only one player defecting may be too severe. On the other hand, maintaining two-way cooperation at A^c also may not be the best response to a one-player defection. For example, in Figure 2, mutual (three-way) defection yields a higher payoff than two-way cooperation at A^c , that is,

$$P_1 (A^d, A^d, A^d) = 12 > 11 = P_1 (A^c, A^c, A^d).$$

Fortunately, players 1 and 2 have other options besides A^d and A^c . They may find it best to choose some other action, somewhere between A^d and A^c , that yields payoffs greater than three-way mutual defection. If they cooperate properly, their (mutual) motivation is to choose an action that maximizes their joint payoff against the defecting third player. We call this action the implicit coalition action, A^{ic} . For a three-player game with player 3 as the defector, it is defined (for player 1) as:

$$P_1 (A^{ic}, A^{ic}, A_3) = \max_{A_j} \{P_1 (A_i, A_i, A_3)\}$$

In the example above, $A^{ic} = 13/9 \approx 1.444$ for any third-player action. In general, the best coalition price will depend on the third player's action, but in the first GPD game it is invariant, just like A^d . In an N -player game there are $N-2$ possible coalition actions corresponding to coalitions of 2, 3, ..., $N-1$ players (One might also wish to define A^d and A^c as coalition prices for coalitions of 1 and N players, respectively.)

² Abreu (1986) has recently proposed a class of strategies known as "carrot and stick" strategies that use severe punishments (as low as A^c and even lower) as a credible threat to enforce maximally collusive behavior.

	$A_1 = A^c$ $A_2 = A^c$	$A_1 = A^c$ $A_2 = A^{ic}$	$A_1 = A^c$ $A_2 = A^d$	$A_1 = A^{ic}$ $A_2 = A^{ic}$	$A_1 = A^{ic}$ $A_2 = A^d$	$A_1 = A^d$ $A_2 = A^d$
$A_3 = A^c$	20.00 20.00 20.00	15.30 15.30 27.20	11.45 11.45 29.25	10.65 22.44 22.44	6.83 18.54 24.46	3.05 20.54 20.54
$A_3 = A^{ic}$	27.20 15.30 15.30	22.44 10.65 22.44	18.54 6.83 24.46	17.72 17.72 17.72	13.85 13.85 19.72	10.01 15.84 15.84
$A_3 = A^d$	29.25 11.45 11.45	24.46 6.83 18.54	20.54 3.05 20.54	19.72 13.85 13.85	15.84 10.01 15.84	11.98 11.98 11.98

Figure 3: Illustration of Payoffs When Actions Are Limited to A^c , A^{ic} , and A^d . (The first, second, and third lines refer to the payoffs to players 1, 2, and 3, respectively.)

Figure 3 illustrates the impact of implicit coalitions. Notice that for a fixed action by player 3, the best cooperative response by players 1 and 2 is always A^{ic} . Furthermore, the subgame between players 1 and 2 is itself a two-player PD in which cooperation becomes A^{ic} while defection is still A^d .

At this point it is clear that there may be some motivation for implicit coalition to form. We have not demonstrated whether or not it is advantageous to play strategies that seek to form coalitions in repeated games. Nonetheless, we propose three strategies for the repeated GPD that recognize and use the concept of implicit coalitions. The first, COALITION, limits action to A^c , A^{ic} , A^d . The second, COALENC, uses the continuous nature of the action set to *encourage coalitions*. The third, GEMERIC, is a generalization of the first two and proves useful when we describe the tournaments. Without loss of generality, we continue to state the algorithms from the perspective of player 1.

COALITION is the simplest possible implicit coalition strategy. It begins each game at A^c . In later rounds it does the following:

$$\text{COALITION: } A_1(t) = \begin{cases} A^c & \text{if } A_2, A_3 \geq A^c \\ A^{ic} & \text{if } \max\{A_2, A_3\} \geq A^{ic} \text{ and } \min\{A_2, A_3\} < A^c \\ A^d & \text{if } A_2, A_3 < A^{ic} \end{cases}$$

where $A_1(t)$ is player 1's action in round t and A_2 and A_3 are actions in round $t-1$.

COALENC is similar to COALITION, except that it recognizes and tries to take advantage of the fully continuous nature of the action set:

$$\text{COALENC: } A_1(t) = \begin{cases} \min\{A_2, A_3, A^c\} & \text{if } A_2, A_3 \geq A^{ic} \\ A^{ic} & \text{if } \max\{A_2, A_3\} \geq A^{ic} \text{ and } \min\{A_2, A_3\} \leq A^{ic} \\ \max\{A_2, A_3, A^d\} & \text{if } A_2, A_3 \leq A^{ic} \end{cases}$$

COALENC will maintain total cooperation (at A^c) if both other players cooperate at A^c , and it will defect to A^d only if both players defect to (or below) A^d . However, at all other times it will hedge toward the implicit coalition price, A^{ic} , by aligning itself with the player closer to A^{ic} .

Finally, we acknowledge that more complex responses are possible in the three ranges of competitors' actions. Indeed, such complex algorithms were entered in the tournaments. We define GENERIC as a generic implicit coalition strategy where f_1 and f_2 are general functions mapping the actions in $t-1$ (or earlier) onto the ranges $[A^d, A^{ic}]$ and $[A^{ic}, A^c]$, respectively:

$$\text{GENERIC: } A_1(t) = \begin{cases} f_2(A_2, A_3) & \text{if } A_2, A_3 \geq A^{ic} \\ A^{ic} & \text{if } \max\{A_2, A_3\} \geq A^{ic} \text{ and } \min\{A_2, A_3\} \leq A^{ic} \\ f_1(A_2, A_3) & \text{if } A_2, A_3 \leq A^{ic} \end{cases}$$

THE TOURNAMENT APPROACH

Given the importance of the PD and its extension, the GPD, it is natural to try to find a "best" strategy for a GPD game that is repeated over many rounds. (In the repeated game, we assume that the payoffs in any one round depend only on the actions in that round, but each player can observe the previous actions by his or her competitors.) Unfortunately, as Axelrod (1981) showed for the classical PD, there is no single best strategy. Against different sets of competitors, different strategies may be best. For example, ALL- A^d is the (unique) best response to a pair of players choosing ALL- A^d , but COALITION is a best response (although not unique) if competitors are known to be playing COALITION.

Axelrod (1980a, 1980b) pioneered a methodology to identify strategies for the classical PD that perform well against a wide range of competitors. In order to generate a rich environment, Axelrod sponsored a contest, inviting game theorists to submit strategies in the form of computer subroutines for a repeated PD game. Each entry "played" every other entry in a round robin tournament. The objective was to earn the highest total score across all games. Entries could be simple or complex; some participants even created strategies that tried to identify opponent's strategies and then act appropriately against them.

The winner was the simplest entry of all, TIT-FOR-TAT which starts cooperatively and in each subsequent period does whatever its opponent did in the previous round. Axelrod described several key properties that helped distinguish the most successful strategies, such as *niceness*, which means "never be the first to defect."

A second tournament was run soon after the first tournament's results were tallied. This time Axelrod received 62 entries from participants, representing a wide range of ages, disciplines, and geographic Origins. The winner, once again, was TIT-FOR-TILT, suggesting that its first-round victory was no fluke. The second tournament reconfirmed the importance of niceness; Axelrod also identified pivotal properties such as *forgiveness* (i.e., do not be too severe when punishing opposing defections), *provocability* (i.e., never let an opposing defection go unacknowledged), and *lack of envy* (i.e., do not intentionally try to reduce competitors' scores). In essence, cooperation can be achieved based upon appropriate reciprocity. Axelrod's tournaments have provoked praise and criticism, but they have raised a number of interesting ideas. We seek to apply the tournament methodology to study implicit coalitions in the GPD.

MITCSI: THE FIRST GPD TOURNAMENT

In November 1984 we announced our first tournament (named, MITCSI for MIT Competitive Strategy Tournament), similar in design to Axelrod's second tournament but featuring the Generalized Prisoner's Dilemma. The game was posed as a managerial problem with price as the sole strategic variable. Each game in the tourney was the repeated GPD game defined in equation 1 with three programmed strategies choosing actions each period from a continuous range. As in Axelrod's tournament, each possible grouping of entries engaged in five repeated games, and the overall winner was the strategy that amassed the highest total score across all games in which it participated. Contestants were given full information about the payoff function used, and in every game each player had access to the past actions of all three players in that game. Entries were submitted in the form of FORTRAN IV subroutines.

By July 1985, we had received over 40 algorithms (including several duplicates) from a diverse group of participants around the world. The field of entrants included economists, political scientists, game theorists, marketing academics, and managers. Several universities and major corporations submitted the best and most creative entries they found after running their own mini tournaments. Thus the pool of algorithms available for this first empirical analysis of the GPD contains a wide variety of creative efforts from some very strategically minded people.

DESCRIPTION OF ENTRIES

Many entrants, having learned Axelrod's lessons, attempted to generalize TIT-FOR-TAT. Six strategies recognized implicit coalitions and incorporated the implicit coalition action, A^{ic} , into their algorithms. We label these algorithms IC for implicit coalition. Several IC entries fit into the GENERIC framework, including some that used very complex functions form and f_1 and f_2 involving many of the previous decisions of each competitor, not just their most recent actions.

Most algorithms tried to incorporate continuous alternatives, but participants did so in a variety of ways, including:

- MIN: Start at A^c . In all subsequent rounds, choose the *minimum* of your competitors' action from the previous rounds.
- MAX: Start at A^c . In all subsequent rounds, choose the *maximum* of your competitors' action from the previous rounds.
- AVG: Start at A^c . In all subsequent rounds, choose the *average* of your competitors' action from the previous rounds.

Each of these strategies leads to very different types of behavior. MAX is extremely forgiving but not highly provokable. (Recall that "higher" action is more cooperative.) Two MAXs, playing together against a nasty competitor, will remain at A^c , ignoring the exploitative moves by the third player. On the other hand, MIN is extremely competitive, and raises its action only if both competitors do so first. AVG is the most moderate of the three,

trying to balance forgiveness and provocability at the same time. (Axelrod himself entered AVG into the tournament.)

A slightly more complicated generalization of TIT-FOR-TAT is MXCM (pronounced “maxcum”) which also starts at A^c , but in later rounds mimics the previous action of the strategy (not including itself) with the greatest cumulative score at that point in the game. Thus MXCM does not attempt to use the previous action of *both* competitors – it considers only the stronger of the two (in terms of cumulative score) and adopts the passive mimicking strategy, just like TIT-FOR-TAT in the two-player game. Like AVG, MXCM is both provokable and forgiving since it follows any action made by the leading strategy. However, MXCM is distinguished from AVG because it is able to ignore ineffective strategies. In contrast, AVG will always give equal weight to the actions of both competitors, regardless of how well they perform.

Other entrants made no attempt to generalize TIT-FOR-TAT. Some used different variants of a strategy suggested by Friedman (1971) that begins each game at A^c and stays there until any competitor defects, in which case it goes to A^d and stays there for all subsequent rounds. We label this type of strategy XTRM.

A few participants chose constant strategies (e.g., always cooperate [ALL- A^c], always defect [ALL- A^d], or always be envious [ALL- A^e]), or random (RND) strategies that chose actions randomly from the range $[A^d, A^e]$ or used a random walk technique. Hence most of the algorithms could be classified into one of eight broad categories: MIN, MAX, AVG, MXCM, IC, XTRM, constant action, or RND.

Beyond these general descriptions of strategy types, the entries differed due to specific tactics or feature that were frequently employed. For example: Following Axelrod (1984), strategies that start the game cooperatively and are never the first to cut action below A^c are termed nice, as opposed to nasty strategies, which can be the first to defect.

Self-awareness allows strategies to consider the previous decisions of all three players (not just the two competitors) when choosing actions. This feature tended to reduce cycling and echo effects.

Many strategies restricted their actions to the range $[A^d, A^c]$, since there is no way to increase payoffs by choosing actions outside of this range. Strategies that were willing to go below A^d or above A^c are known as *unbounded*.

Some strategies tried to induce cooperation by raising actions slightly above the level specified by a general strategy type (e.g. play MIN but add a few units to the minimum action). These strategies have *action-raising initiative*. On the other hand, some strategies have *action-cutting initiative*, that is, they were willing to go below the specified action level, usually in an attempt to punish an earlier cut made by a competitor.

Finally, several strategies occasionally used the *envious* action A^e to try to outscore their competitors (rather than maximizing their own scores).

These descriptions are admittedly vague. The actual implementation of some of these features can vary greatly from strategy to strategy. For instance, action-raising algorithms can vary the magnitude and frequency of their increases. For ease of exposition, we ignore fine-grained differences among algorithms, since the mere presence of a particular feature was generally more important than the manner in which it was implemented.

RESULTS AND INTERPRETATION

Table 1 presents a summary of the strategies and their performances. The strategies are ranked by their average score per round. For comparison, mutual cooperation pays 20 units per period to each player, while each period of mutual defection (i.e., all three players choosing A^d) yields approximately 12 units to each player. (See Figure 3 for additional payoff comparisons.)

The winning algorithm, entered by Terry Elrod of Vanderbilt University, was the simplest possible IC strategy, COALITION. Recognizing implicit coalitions proved to be the single most important factor in the tournament. The top four algorithms in the tournament recognized the coalition properly, and all six IC strategies finished in the top 10 overall. Another interesting factor was how the strategies dealt with the continuous nature of the actions. Most entrants used simple heuristics (e.g. MIN, MAX, AVG) to address this problem, with varying degrees of success. The standard averaging strategy (i.e. nice, bounded AVG with no self-awareness, envy, or action-raising/-cutting initiative) finished in sixth place, easily beating standard MAX (ranked eleventh) and standard MIN (twelfth).

TABLE 1
Official MITCS1 Results

Rank	Entrant	Strategy Type	Features ¹					Average Score per Round
1	Terry Elrod	IC						17.182
2	(anonymous)	IC	S	U	R		E	17.172
3	Avraham Beja & Shlomo Kalish	IC	S				C	17.172
4	Steve Shugan	IC						17.157
5	(MIT) ²	AVG ³	S					17.157
6	Gary A. Lines	AVG						17.104
7	(MIT)	MIN	S		R			17.063
8	Steve Shugan	IC	S					17.014
9	Beja & Kalish	IC	S				C	16.927
10	(MIT)	MXCM	S					16.914
11	Terry Elrod, John Roberts	MAX						16.879
12	Gary Gaeth & Gerard Tellis, Terry Elrod, Gary A.Lines	MIN						16.851
13	James M. Lattin	MAX	S		R			16.830
14	John A.Cadley	XTRM			R			16.720
15	Steve Borgatti	MIN	S					16.682
16	Steve Borgatti	MIN					C	16.523
17	John A. Cadley	XTRM						16.519
18	(MIT)	⁴		U				16.389
19	RobertbAxelrod	AVG		U				16.335
20	John Roberts	AVG ⁵		U				16.305
21	Barbara Bruner & James Olver	AVG	N	S			C	16.127
22	Robert F. Bordley	MIN		U				15.976
23	Robert E. Marks	XTRM ⁶		U			E	14.535
24	Robert E. Marks	XTRM ⁶		U			E	14.497
25	James M. Lattin	ALL-A ^c						14.351
26	Shlomo Maital	MXCM	N	U				13.944
27	Beja & Kalish	RND ²	N	S				13.809
28	Robert E. Marks	⁸	N	U			E	13.763
29	Steve Borgatti	MIN	N	S			C	13.740
30	Roland Rust, Robert F. Bordley,, John Roberts	ALL-A ^d	N					13.637
31	Beja & Kalish	RND ⁷	N	S				13.575
32	(MIT)	RND ⁹	N					13.447
33	Shlomo Maital	MXCM	N	U				13.384
34	Robert F. Bordley	MIN	N	U				12.918
35	Kenneth L. Stott, Jr.Francis J. Vasko & Floyd E. Wolf	ALL-A ^c	N	S	U	R	E	12.151
36	Robert E. Marks	ALL-A ^c	N	U			E	9.909
37	(anonymous)	MIN ¹⁰	N	U	R			9.684
38	(anonymous)	AVG ¹¹		U	R	C		9.643

(continued)

TABLE 1 Continued

NOTES:

- (1) Default features include niceness, no self-awareness, bounded actions, no action-raising or cutting initiative and no envy. Exceptions are noted as N for nastiness, S for self-awareness, U for unbounded actions, R for action-raising initiative, C for action-cutting initiative, and E for envy.
- (2) (MIT) denotes an algorithm entered by a member of the MIT community that was not eligible to win the tournament. Post-tournament testing indicates that the inclusion of these entries does not affect the ordering of the top algorithms.
- (3) Weighted average of all three players, actions, using cumulative scores as weights.
- (4) Mimics previous move of one opponent on odd rounds. other opponent on even rounds.
- (5) Geometric mean of opponents' previous actions.
- (6) Stays at A^c for two opposing defections before going to A^e .
- (7) Random walk centered around $\frac{1}{2} (A^c + A^d)$.
- (8) Mimics actions of one opponent chosen at start of game: actions limited to range(A^c , A^{lc})
- (9) Uniform random variable between A^d and A^c .
- (10) Only algorithm to introduce action increase above A^c .
- (11) Only algorithm to introduce action increase cuts below A^c .

Several of the descriptive features were highly influential. First and foremost is niceness, reconfirming the findings of Axelrod. Nice algorithms were able to reap great benefits by avoiding the short-term temptation to defect. The best nasty strategy played standard AVG most of the time, but would occasionally make small cuts as long as both competitors remained at A^c . If either competitor responded to these cuts, this strategy would return to standard AVG for the remainder of the game, this clever form of exploitation helped make this algorithm far more successful than other nasty entries, but still could not provide any better than a twenty-first place finish.

Other important features were boundedness and lack of envy. Only one successful algorithm ever exhibited envious behavior, but that strategy (ranked second) would only go to A^c if both competitors were at or below A^e in the previous round, a fairly rare occurrence. Perhaps if this second-ranked strategy did not try to battle envious competitors on their terms, it might have been able to win the tournament.

The value of bounded actions can be seen by comparing standard AVG and MIN (ranked sixth and twelfth, respectively) to their equivalent but unbounded counterparts (ranked nineteenth and twenty second, respectively). Boundedness was worth nearly 0.70 units per round to AVG and nearly 0.90 units per round to MIN,

Several entrants found self-awareness to be a blessing. For example, some strategies, unlike TIT-FOR-TAT, considered their own previous decisions in determining future actions, for example, averaging across all three players (ranked fifth), and three-player MXCM (tenth place). But self-awareness was a curse to others, including those who used it as a ratchet on actions. The algorithms ranked eighth and fifteenth, for example, only let their actions move downwards, regardless of the cooperative gestures made by their competitors.

Little can be said about the effectiveness of action-raising and action-cutting initiative. Some of the nice, bounded entries were able to encourage cooperation and discourage cheating with appropriate rewards and penalties, but these successes were counterbalanced by the unsuccessful strategies that brought on their own demise by raising or cutting actions too much at the wrong times.

TABLE 2
Revised MITCS1 Results

<i>New Rank</i>	<i>Entrant</i>	<i>Old Rank</i>	<i>Strategy Type</i>	<i>Features¹</i>	<i>Average Profits per Round</i>
1	COALENC	-	IC		17.353
2	Terry Elrod	1	IC		17.257
3	(anonymous)	2	IC	S U R E	17.249
4	Avraham Beja & Shlomo Kalish	3	IC	S C	17.245
5	Steve Shugan	4	IC		17.235
6	(MIT) ²	5	AVG ³	S	17.225
7	Gary A. Lines	6	AVG		17.175
8	(MIT)	7	MIN	S R	17.136
9	Steve Shugan	8	IC	S	17.101
10	Beja & Kalish	9	IC	S C	17.000

NOTES:

- (1) Default features include niceness, no self-awareness, bounded actions, no action-raising or -cutting initiative, and no envy. Exceptions are noted as N for nastiness, S for self-awareness, U for unbounded actions, R for action-raising initiative, C for action-cutting initiative, and E for envy.
- (2) (MIT) denotes an algorithm entered by a member of the MIT community.
- (3) Weighted average of all three players' action, using cumulative scores as weights.

Table 1 seems to depict a tight three-way battle for first place. However, it should be noted that each algorithm played in nearly 1,000 three-player matches in each of the five games in the tournament. This information, combined with the fact that each match lasted approximately 200 rounds, implies that each strategy chose actions in nearly 1 million total rounds. Thus a difference of .01 units on a payoff-per-round basis is equivalent to a 10,000 units difference in total score.

AN ALTERNATIVE CHAMPION

The winning algorithm, COALITION, was the only highly ranked strategy that did not acknowledge the continuity of actions. Apparently, none of its top rivals could use the continuous nature of the action space enough to overcome the winner's discrete simplicity. However, this does not imply that the task is impossible; COALENC, described earlier, would have easily won the tournament had it been entered.

Table 2 shows the top 10 entries in the revised tournament with COALENC included. Note that the relative rankings of the original strategies are unchanged, although average scores have increased because of the presence of the cooperative newcomer. The margin of victory for the new algorithm is quite significant; the gap between first and second place is larger than the margin between second and seventh place.

More importantly, the success of this algorithm is not very sensitive to variations in the competitive environment. Extreme changes, such as doubling the presence of all nasty entries, usually cannot unseat this new winner. Many of the procedures that Axelrod used to demonstrate the robustness of TIT-FOR-TAT have been applied to this tournament, with strong results favoring COALENC.

MITCS2: THE SECOND GPD TOURNAMENT

One of the unique aspects of the payoff function used in MITCS1 is separability, which leads to unique, invariant values for A^d and A^{ic} . Because the implicit coalition action never changes, it is relatively easy for coalition-seeking algorithms to achieve their goal. In more general situations, the best action for a coalition should depend on the actions of noncoalition players. For example, the coalition response to an envious player might be harsher (i.e., lower coalition action) than the coalition response to a small defection. (Recall once again that higher action is defined as more cooperative). With this in mind, we sought to determine whether the success of COAUTION and COALENC was unique to the payoff function and competitive environment of MITCS1, or whether it could be replicated in an environment that is potentially less favorable to implicit coalitions.

Soon after we completed the analysis of MITCS1, we announced a second tournament, MITCS2, with the following payoff function:

$$\Pi_1 = 200(8 - 6A_1 + A_2 + A_3)(A_1 - 1) - 180 \quad (2)$$

Equation 2 corresponds to a linear demand function in economics. As before, payoffs are symmetric and the equation satisfies the GPD conditions. The scaling constants were chosen to match closely the payoffs in MITCS1; full cooperation ($A_1 = A_2 = A_3 = A^c$) pays 20 units per player per round, and full defection ($A_1 = A_2 = A_3 = A^d$) pays 12 units per player per round.

The key difference between MITCS1 and MITCS2 is that the short-term payoff-maximizing action (A^d) and the implicit coalition action (A^{ic}) now depend on competitors' actions. Specifically, for fixed competitive actions,

$$A_1^d = \frac{14 + A_2 + A_3}{12} \quad (3)$$

$$A_1^{ic} = A_2^{ic} = \frac{13 + A_3}{10} \quad (4)$$

For example, if player 3 chooses ALL-1.40, then $A_1^{ic} = A_2^{ic} = 1.440$. However, if player 3 chooses ALL- A^d via equation 3, then $A_3^d \approx 1.407$, and $A_1^{ic} \approx A_2^{ic} \approx 1.441$. All entrants were aware of MITCS1 and the success of COALITION and COALENC. Each subset of three entries was matched for five games of 200 rounds³ and the winner was the strategy with the highest total (or average) payoffs.

TOURNAMENTS RESULTS

By fall of 1986, 32 entries had been submitted to MITCS2. Five strategies were thrown out due to coding errors or illegal tactics. The remaining 27 entries were combined with 11 strategies carried over (some with slight modifications) from MITCS1. These strategies were

³ Since no entries used any explicit end-game maneuvers, the game length was fixed at 200 rounds For all games.

included again because they led to interesting pricing behavior in the original tournament. Finally, suggestions from other individuals who did not wish officially to participate led to 6 more submissions, thus rounding out the field of 44 unique entries.

A brief description of each entry is shown in Table 3, where the entries are ranked by average scores per round.

TABLE 3
MITCS2 Official Results

<i>Rank</i>	<i>Entrant</i>	<i>Strategy Type¹</i>	<i>N= Nasty</i>	<i>Lower Bound</i>	<i>Mean Score per Round</i>	<i>Description²</i>
1	Robert E. Marks	COALENC		A^d	17.097	$A^{ic} = (26 + A_2 + A_3) / 20$
2	Robert L. Bishop, Tony Haig	COALENC		1.4	17.096	Original COALENC (with $A^{ic} = 13/9$)
3	Paul R. Pudaite	COALENC		1.4	17.091	$A^{ic} = (13 + A_3)/10$; looks back two rounds
4	John Hulland	COALENC		1.4	17.085	$A^{ic} = (13 + A_3)/10$
5	Neil Bermann	COALENC		A^d	17.084	$A^{ic} = (13 + A_3)/10$
6	Tony Haig	COALENC		A^d	17.075	$A^{ic} = (13 + A_3)/10$; looks back two rounds
7	James M. Lattin	COALENC		A^d	17.065	$A^{ic} = 1.44$
8	(MITCS1 # 7) ³	MIN.R		A^d	17.064	Standard MIN with random 2c price increases
9	Scott A. Neslin	AVG.S		1.4	17.063	Linear leaning model; complex averaging procedure
10	Scott A. Neslin	AVG.S		1.4	17.042	Variation of #9 above (i.e., different parameters)
11	- ⁴	MXCM.S		A^d	17.025	Mimics previous price of <i>second-best</i> firm
12	Robert L. Bishop	COALENC		A^d	16.993	$A^{ic} = 13/9$; uses $\max\{A_2, (193 + A_3)/10\}$ when $A_2 < A^{ic}$
13	----	AVG		A^d	16.989	Gradually shifts from MAX to MIN as game progresses
14	(MITCS1 # 5)	AVG.S		A^d	16.970	Weighted average of all 3 players' previous prices
15	-	AVG.S		A^d	16.968	Unweighted average of all 3 players' previous 3 prices
16	Karel Najman	AVG.S			16.964	Unweighted average of all 3 players' previous prices
17	(MITCS1 # 6)	AVG		A^d	16.932	Standard AVG; average of opponents previous prices
18	James M. Lattin	-		A^d	16.932	Complex adaptive learning model
19	(MITCS1 # 11)	MAX		1.4	16.926	Standard MAX; maximum of opponents' previous prices
20	Karel Najman	AVG		-	16.908	Unbounded AVG
21	Terry Elrod	COALITION		1.4	16.907	Same as official winner of MITCS1 but with $A^{ic} = 85/59$
22	James M. Lattin	MIN		1.4	16.887	Plays AVG in round 2, MIN thereafter
23	(MITCS1 # 21)	COALENC	N	A^d	16.765	Modified version of top nasty entry in MITCS1
24	Neil Bermann	COALITION		A^d	16.754	COALITION with varying A^{ic} ; $A^{ic} = (13 + A_3)/10$
25	Chris Jones	MIN		-	16.707	Standard MIN; minimum of opponents' previous prices
26	(MITCS1 # 10)	MXCM.S		A^d	16.679	Mimics previous price of best (most profitable) firm
27	Karel Najman	MIN		1.4	16.567	Unbounded MIN
28	(MITCS1 # 17)	XTRM		A^d	16.187	Play A^c until anyone cuts price; play A^d thereafter

(continued)

Rank	Entrant	Strategy Type ¹	N= Nasty	Lower Bound	Mean Score per Round	Description ²
29	-	AVG	N	A^d	16.162	Start at 1.4 then play AVG
30	-	AVG	N	1.4	15.992	Weighted AVG with random weights
31	James M. Lattin	MAX	N	1.4	15.776	Start at 1.4 then play MAX
32	Karel Najman	ALL-1.44	N	1.4	15.719	Always choose 1.44
33	(MITCS1 # 26)	MXCM	N	A^d	15.264	Start at 1.4 then mimic previous price of best opponent
34	-	XTRM		1.4	15.226	Choose 1.4 for 2 rounds after a price cut, then return to A^c
35	(MITCS1 # 35)	ALL- A^d .R	N	1.4	14.512	Raise price above A^d if profits exceed opponents' profits
36	Peter J. Brock	AVG	N	1.4	14.126	Start at 1.4 then choose average of A^d and $(A_2 = A_3) / 2$
37	Paul R. Pudaite	-	N	-	14.033	Choose $A_1 = (6 + A_2 + A_3) / 6$ to maximum joint profits
38	Karel Najman	ALL- A^d	N	A^d	13.979	Start at 1.5; thereafter choose A^d
39	Paul R. Pudaite	ALL- A^d	N	A^d	13.854	Start at 17/12; thereafter choose A^d
40	Peter J. Brock	ALL- A^d	N	A^d	13.734	Start at 1.4; thereafter choose A^d
41	James M. Lattin	ALL-1.39	N	1.39	13.213	Always choose 1.39
42	Peter J. Brock	ALL- A^d .C	N	--	13.089	Choose $A_1 = (7 - A_2 - A_3) / 6$ to hurt non-ALL- A^d players
43	(MITCS1 # 32)	RANDOM	N	1.333	12.817	Uniform random variable between 1.333 and 1.5
44	(MITCS1 # 36)	ALL- A^d	N	-	11.754	Act enviously, i.e., maximize <i>share</i> of industry

NOTES:

- (1) The basic strategy types are defined in the descriptions and in the text. The suffix "S" refers to each strategy with *self-awareness*. the "R" refers to strategies with *action-cutting initiative*.
- (2) Player 3 is assumed to be less cooperative of players 2 and 3.
- (3) This denotes a strategy that was an official entry in MITCS1. The number refers to its ranking in that tournament.
- (4) This denotes a strategy based on an informal suggestion

Table 3 shows two striking patterns. Most of the COALENC generalizations cluster toward the top, and 15 out of the 16 bottom entries are nasty (i.e., willing to initiate defections). One pattern that is not immediately obvious, however, is the possible link between the success of the COALENC entries and the method of choosing a value of A^{ic} .

The winning entry, submitted by Robert Marks of the Australian Graduate School of Management, features an unusual type of coalition action. It uses equation 4 to calculate an A^{ic} against player 3 and averages this action with an A^{ic} calculated against player 2. Thus, for instance, if $A_2 = 1.50$ and $A_3 = 1.40$, this algorithm would act like COALENC with $A^{ic} = (1.44 * 1.45) / 2 = 1.445$, as compared to a A^{ic} of 1.44 that equation 4 would suggest (and most COALENC entries would use).

At first glance this may seem like an inefficient rule, since it will often lead to coalitions with an action slightly above the "optimal" A^{ic} . But notice which routine came in a close second: the original version of COALENC with $A^{ic} = 13 / 9 = 1.444...$. This is also a relatively high (i.e., more cooperative) coalition action; it will exceed the A^{ic} suggested by equation 4 whenever the noncooperative player is below 1.444.... A pattern emerges: The top two strategies consistently choose higher coalition actions than any of the other COALENC entries. As further evidence, note that the "worst of the best," entry 7, will generally choose the lowest A^{ic} , 1.440.

We briefly summarize some of the other results of interest. First, notice the rather mediocre performance of the entries that attempt to generalize COALITION, as compared to its sterling performance in MITCS1. Part of this drop can be attributed to the different mix of strategies in MITCS2 compared to MITCS1: With the presence of more sophisticated entries (such as the COALENC generalizations), the discrete pricing policy begins to hurt COALITION. This is particularly true when action-cutting exists at moderate levels. But much of COALITION's drop is due to the new payoff function: Without a fixed A^{ic} to rely on, any coalition seeker must be more flexible and forgiving in trying to establish a successful coalition.

Another prominent result from MITCS 1 was the need for a lower bound on one's actions. Most entrants to MITCS2 recognized this idea and used one of two lower bounds—fixed at 1.40 or floating (A^d). The results in Table 3 show no significant advantage for one method or the other. For example, entries 4 and 5 are exactly the same except for their lower bounds, and in each of the five constituent games in the tourney, these entries finish with nearly identical scores. This finding should not be considered too surprising; after all, when action cutting is severe enough to require bounded actions, A^d is usually quite close to 1.40 anyway.

Finally, another result worth mentioning is the relative performance of three standard algorithms. In MITCS2, just as in MITCS1, AVG (entry 17) earns higher payoffs than MAX (entry 19), and both beat out MIN (entry 25). The value of having bounded actions can be seen once again by comparing entries 17 to 20 and 25 to 27. Boundedness does not appear to be as valuable as in MITCS1, but this is only because of the smaller number of extreme action cutters. Only three entries (41, 43, and 44) ever initiate cuts below 1.40.

A NEW ALTERNATIVE CHAMPION

MITCS2 confirms the importance of the implicit coalition phenomenon but suggests that algorithms can be fine-tuned to achieve greater payoffs. In fact, the higher the target coalition action, the better the algorithm seems to perform. We call this new property magnanimity. The success of the magnanimous entries seems to result from the fact that a high A^{ic} is less likely to be viewed as a noncooperative action. In contrast, a less magnanimous algorithm (e. g., entry 7) often will be mistaken for a defector. Matchups between entry 7 and discrete COALITION strategies with higher A^{ic} 's will quickly degenerate into (A^d, A^d, A^d) behavior because the two potential cooperators cannot agree on a coalition action. Of course, there is a limit to magnanimity; too high a coalition action allows an algorithm to be exploited.

To generate a slightly more magnanimous strategy, we included another potentially beneficial property, self-awareness, into the A^{ic} calculation. If cooperative players incorporate their own previous actions in determining which A^{ic} to choose, the resulting coalition action will tend to be higher and more stable. (It is higher whenever $A_1 > (A_2 > A_3)/2$.) Furthermore, a common A^{ic} calculation would avoid the possibility of different subsets of players seeking different coalition actions. We believe that these two effects would cause environments with general payoff functions to become more like the MITCS1 world, where stable coalitions are easily established and maintained.

TABLE 4
Revised MITCS2 Results

Rank	Entrant	Strategy Type ¹	N= Nasty	Lower Bound	Mean Score per Round	Description ²
1A	CEAVG3	COALENC		A^d	17.170	$A^{ic} = (39 + A_1 + A_2 + A_3) / 30$
1	Robert E. Marks	COALENC		A^d	17.163	$A^{ic} = (26 + A_2 + A_3) / 20$
2	Robert L. Bishop, Tony Haig	COALENC		1.4	17.160	Original COALENC (with $A^{ic} = 13/9$)
3	Paul Pudiate	COALENC		1.4	17.156	$A^{ic} = (13 + A_3) / 10$; looks back two rounds
4	John Hulland	COALENC		1.4	17.149	$A^{ic} = (13 + A_3) / 10$

Our new strategy, named CEAVG3 (for coalition encourager, based on the average of all 3 coalition actions), is still a COALENC strategy; only the coalition reaction function is different. In CEAVG3, instead of calculating and averaging our A^{ic} against players 2 and 3, we perform the same task with respect to *all three players*. The new coalition reaction function, therefore, is

$$A_1^{ic} = \frac{39 + A_1 + A_2 + A_3}{30} \quad (5)$$

Although the coalition actions and payoffs for the new strategy are only slightly higher than those of entry 1, this small increase combined with the moderating influence of the lagged A_1 term helps the new strategy to achieve a first-place finish when placed among the MITCS2 entries. Table 4 shows the revised payoff figures. (Only the top five entries are shown; the overall standings are barely affected by the presence of the new strategy.)

Since CEAVG3 is a COALENC strategy, its behavior (and payoffs) will often be indistinguishable from the other COALENCs. However, in the cases in which these entries do differ, CEAVG3 does well enough to win the revised MITCS2 tournament by a relatively comfortable margin.

TEST OF ROBUSTNESS

No tournament can tell us which single strategy is truly “best,” or which set of strategies will do well in the widest set of environments. But a series of tournaments coupled with some reasoning can raise some valid hypotheses and insights.

COALENC strategies did well in both MITCS1 and MITCS2, but these tournaments represent relatively “nice” environments. To determine the sensitivity of COALENC strategies to environments, we performed a test of robustness, similar to the post-tournament work of Axelrod (1984). We generated 200 new environments using different combinations of the MITCS2 entries. We first used a stepwise procedure to identify a subset of eight representative entries that faithfully reproduce the overall payoffs and standings of MITCS2, using only a small fraction of the full tournament. The eight representatives (7, 16, 20, 23, 28, 33, 37, and 41) form an environment involving 36 games with each of the MITCS2 entries but yield overall average payoffs that have a correlation coefficient of 99.49% with the scores from the full tournament (5, 175 games per entry).

TABLE 5
Strategy Performance in Simulated Environments

<i>entry</i>	<i>score</i>	<i>entry</i>	<i>score</i>	<i>entry</i>	<i>score</i>	<i>entry</i>	<i>score</i>	<i>entry</i>	<i>score</i>
11	16.186	1A	16.866	1A	17.364	1A	17.833	12	18.867
1A	16.162	2	16.861	2	17.357	2	17.826	21	18.821
2	16.159	1	16.860	1	17.356	1	17.825	1A	18.780
1	16.155	3	16.859	3	17.352	3	17.822	2	18.774
3	16.150	4	16.853	9	17.351	5	17.817	1	18.774
4	16.144	5	16.852	4	17.346	4	17.816	11	18.771
5	16.143	6	16.841	5	17.346	8	17.809	3	18.766
6	16.132	7	16.839	10	17.343	6	17.806	5	18.762
7	16.106	8	16.819	8	17.337	9	17.804	4	18.761
9	16.105	9	16.818	-6	17.335	7	17.800	24	18.757

←	●	●	●	●	●	→
<i>Nastiest</i>		<i>Moderately Nasty</i>	<i>Mid - range</i>	<i>Moderately Nice</i>	<i>Nicest</i>	

<i>entry</i>	<i>type</i>	<i>entry</i>	<i>type</i>	<i>entry</i>	<i>type</i>
1A	COALENC	5	COALENC	10	AVG. S
1	COALENC	6	COALENC	11	MXCM. S
2	COALENC	7	COALENC	12	COALENC
3	COALENC	8	MIN. R	21	COALITION
4	COALENC	9	AVG. S	24	COALITION

To generate each stimulated environment, we took random combinations of each of the eight representatives, also accounting for the residuals between the actual and mini tournament payoffs. This procedure was repeated 200 times, thereby producing a wide range of environments.

As a proxy for the niceness or nastiness of each environment, we use the average payoffs across all 45 entries. The 200 environments are sorted by this index and broken into five equal-sized groups. We ranked the score for each strategy within each group. Table 5 summarizes the results by giving the top 10 finishers in each environment. For ease of reference, the basic strategy types are shown below for each listed entry. (The entry numbers refer to MITCS2 rankings.)

Table 5 shows a clear, consistent pattern supporting the results of Tables 3 and 4. The top strategies are very stable in moderate environments and fall only slightly in more extreme cases. It is encouraging to see that the COALENC entries perform so well even in very nasty environments. Even in the single nastiest environment, where over 60% of the random weight is allocated to the nasty representatives, four COALENC entries finish in the top 10, and only one nasty strategy finishes in the top 25.

One surprise that emerged out of the simulations is entry 11. This strategy is based on a very unusual notion: It identifies the second-best player in each game (in terms of cumulative payoffs) and mimics that player's previous action. This rule adapts very well to extreme environments (good or bad) since it goes along with coalitions in a most magnanimous way (good in nice environments) but never initiates coalition behavior (good in nasty environments). If we look at alternative measures of performance, such as number of first-

place finishes in the 200 simulations, then 11 appears to be even stronger. It is the winner in 50 of the environments, more than any other MITCS2 entry.⁴

SUMMARY

This article has examined the role of implicit coalitions in a generalized prisoner's dilemma. We find the GPD interesting because it extends the classical PD to more realistic situations of more than two (but not many) players and it gives players the option of choosing actions from a continuous set. When we extend the PD to the GPD we find the possibility of implicit coalitions, that is, coalitions of cooperating players in an otherwise unfriendly world. We also expect, intuitively, that strategies that use the continuous nature of the action space will do better than those that do not.

We tested our conjectures in two three-player GPD's, as described by equations 1 and 2. Our methodology was that of computer tournaments. In both tournaments, implicit coalitions proved to be the key feature that distinguished the most successful strategies (in terms of average score). In MITCS1 a simple discrete COALITION strategy won and other coalition strategies fared well. However, a specific coalition-encouraging strategy, COALENC, would have won had it been entered. In MITCS2, several different variants of COALENC did surprisingly well, especially considering the differences in the payoff function from MITCS1 to MITCS2. Among the COALENC entries, magnanimity seemed to distinguish the very best algorithms. Finally, COALENC strategies held up very well in a variety of hypothetical environments, although at least one alternative algorithm did well in the nastiest of environments. Despite the fine showing of CEAVG3 and the other COALENC entries, we dare not make our claims too strong. The GPD is a rich and complex problem and our tournaments only begin to tap its complexity. Nonetheless, we do feel confident that implicit coalitions are important and should be considered in any situation modeled by a GPD.

As in all research, interesting questions remain. Beyond the obvious questions of more than three players, alternative payoff functions, and still more complex algorithms, we feel that further investigation of magnanimity and further exploitation of continuous action are warranted. Algorithms that have greater adaptability to recognize competitors deserve attention.

One interesting theme that emerged was that it often pays to be more cooperative than a simple one-for-one matching policy (such as TIT-FOR-TAT) would suggest. For example, in the second tournament, magnanimity implies that strategies should shade toward being more cooperative when choosing implicit coalition actions. Even in COALENC, when both actions are below A^{ic} , the strategy chooses the more cooperative action of the other two players. Further investigation of this theme should prove fruitful.

Beyond computer tournaments, there are possibilities for GPD experiments on human subjects and descriptive research to determine which real-world conflict situations are best modeled by OPDs and implicit coalitions. Finally, we view implicit coalitions as an excellent

⁴ Although entry 11 is most adept at winning, it does have its bad moments. It finishes out of the top ten 40.5% of the time, including a low of twenty-ninth place in one environment. CEAVG3, for comparison, is far more robust with only 15.5% of its ranking below the top ten, never lower than fifteenth place.

concept to examine the overlap (or differences) in the approaches used by cooperative and noncooperative game theory to study multiple player conflict situations.⁵

REFERENCES

- ABREU, D. (1986) "Extremal equilibria of oligopolistic super games." *Journal of Economic Theory* 39: 191 -225.
- AXELROD, R. (1980a) "Effective choice in the Prisoner's Dilemma." *J. of Conflict Resolution* 24 (March): 3-25.
- AXELROD, R. (1980b) "More effective choice in the Prisoners Dilemma" *J. of Conflict Resolution* 24 (September): 379-403.
- AXELROD, R. (1981) "The emergence of cooperation among egoists." *Amer. Pol. Sci. Rev.* 75: 306-318.
- AXELROD, R. (1984) *The Evolution of Cooperation*. New York: Basic Books.
- BARNABY, F. (1987) "The nuclear arsenal in the Middle East." *Technology Rev.* 90 (May-June): 17-35.
- DAWES, R. M. (1980) "Social dilemmas." *Annual Rev. of Psychology* 31: 169-193.
- FRIEDMAN, J. W. (1971) "A noncooperative equilibrium for super games." *Rev. of Economic Studies* 38:1-12.
- GOEHRING, D. J. and J. P. KAHAN (1976) "The uniform n-person Prisoners Dilemma game." *J. of Conflict Resolution* 20 (March): 111-128.
- GRIFFIN, A. (1987) "Cooperative R&D between competitors." Working Paper, MIT, Sloan School of Management.
- HAMBURGER, H. (1979) *Games as Models of Social Phenomena*. San Francisco: W. H.
- HARDIN, (1968) "The Tragedy of the Commons." *Science* 162: 1243-1248.
- SCHELLING, T. C. (1973) "Hockey helmets, concealed weapons, and daylight saving: a study of binary choices with externalities." *J. of Conflict Resolution* 17: 381-428.
- SHUBIK, M. (1959) *Strategy and Market Structure*. New York: John Wiley.
- TAYLOR, M. (1976) *Anarchy and Cooperation*. New York: John Wiley.

⁵ The competitive strategy tournaments were supported by the Marketing Center at MIT, with assistance from the Marketing Science Institute in Cambridge, MA.