Cooperative Equilibrium

(Extended Abstract)

Joseph Y. Halpern
Computer Science Department
Cornell University
Ithaca, NY 14853
halpern@cs.cornell.edu

Nan Rong
Computer Science Department
Cornell Universtiy
Ithaca, NY 14853
rongnan@cs.cornell.edu

ABSTRACT

We propose a new equilibrium concept, *perfect cooperative equilibrium (PCE)*, which may help explain players' behavior in games where cooperation is observed in practice. We also consider a few related equilibrium concepts that take into account the degree of cooperation.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multiagent systems*

General Terms

Economics, Theory

Keywords

Cooperative equilibrium, Nash equilibrium, game theory, cooperation, punishment

1. INTRODUCTION

Nash equilibrium (NE) is perhaps the most widely used solution concept in game theory. However, there are a number of games where NE does a poor job in predicting behavior. For example, consider Prisoner's Dilemma, in which two prisoners can choose either to defect or cooperate, with payoffs as shown in the table below. Although the only NE is (Defect, Defect), people often play (Cooperate,Cooperate).

	Cooperate	Defect
Cooperate	(3,3)	(0,5)
Defect	(5,0)	(1,1)

For another example, consider the Traveler's Dilemma [1]. In this game, two travelers have identical luggage that is damaged (in an identical way) by an airline. The airline offers to recompense them for their luggage. They may ask for any dollar amount between \$2 and \$100. There is only one catch. If they ask for the same amount, then that is what they will both receive. However, if they ask for different amounts—say one asks for $\$m$ and the other for $\$m'$, with $m < m'$—then whoever asks for $\$m$ (the lower amount) will get $\$(m+2)$, while the other traveler will get $\$(m-2)$. A little calculation shows that the only NE in the Traveler's Dilemma is $(2,2)$. Nevertheless, in practice, people (even game theorists!) do not play (2,2). Indeed, when Becker, Carter, and Naeve [2] asked members of the Game Theory Society to submit strategies for the game, 37 out of 51 people submitted a strategy of 90 or higher among which 10 sumbitted the cooperative strategy 100 – which is also the most popular strategy; the winning strategy (in pairwise matchups against all submitted strategies) was 97. Only 3 of 51 people submitted the "recommended" strategy 2. In this case, NE is neither predictive nor normative; it is far from what people play, and it produces quite poor results.

In both of these games, people exhibit behavior that can be viewed as cooperative, which cannot be explained by NE. Are there rules underlying cooperative behavior? In this paper, we propose a new equilibrium concept: *perfect cooperative equilibrium* (PCE). PCE may help explain players' behavior in games where cooperation is observed in practice. A player's payoff in a PCE is at least as high as in any NE. However, a PCE does not always exist. We thus consider α-PCE, where α takes into account the degree of cooperation; a PCE is a 0-PCE. Every game has a Pareto optimal *max-perfect cooperative equilibrium (M-PCE)*; that is, an α-PCE for a maximum α. We show that M-PCE does well at predicting behavior in quite a few games of interest.

2. PCE

We now introduce PCE. For ease of exposition, we focus here on finite *normal-form* games $G = (N, A, u)$, where $N = \{1, \ldots, n\}$ is a finite set of players, $A = A_1 \times \ldots \times A_n$, A_i is a finite set of possible actions for player i, $u = (u_1, \ldots, u_n)$, and u_i is player i's utility function. Players are allowed to randomize. A strategy for player i is thus a distribution over actions in A_i; let S_i represent the set of player i's strategies. Let $U_i(s_1, \ldots, s_n)$ denote player i's expected utility if the strategy profile $s = (s_1, \ldots, s_n)$ is played. Given a profile $x = (x_1, \ldots, x_n)$, let x_{-i} denote the tuple consisting of all values x_j for $j \neq i$.

Definition 1. *Given a game G, a strategy s_i for player i in G is a best response to a strategy s_{-i} for the players in $N - \{i\}$ if $U_i(s_i, s_{-i}) = \sup_{s'_i \in S_i} U_i(s'_i, s_{-i})$. Let $BR_i^G(s_{-i})$ be the set of best responses to s_{-i} in game G.*

We first define PCE for 2-player games.

Definition 2. *Given a 2-player game G, let BU_i^G denote the best utility that player i can obtain if the other player j*

Cite as: Cooperative Equilibrium (Extended Abstract), Joseph Y. Halpern and Nan Rong, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 1465-1466

best responds; that is,

$$BU_i^G = \sup_{\{s_i \in S_i, s_j \in BR^G(s_i)\}} U_i(s).$$

Definition 3. *A strategy profile s is a perfect cooperative equilibrium (PCE) in a 2-player game G if for all $i \in \{1,2\}$, we have $U_i(s) \geq BU_i^G$.*

Intuitively, in a 2-player game, a strategy profile is a PCE if each player does at least as well as she would if the other player were best-responding. In Prisoner's Dilemma, both (Cooperate, Cooperate) and (Defect, Defect) are PCE.

PCE has a number of attractive properties: every player does at least as well in a PCE as in a NE; every dominant strategy profile is a PCE; and a strategy profile that Pareto dominates a PCE is a PCE.

Despite of these properties, why should or do players play (their part of) a PCE? Consider (one of) the intuitions for NE: players have learned other players' strategies through playing the game repeatedly; they thus best respond to what they have learned. A NE is a stable point of this process: every players' strategy is already a best response to what the other players are doing. This intuition focuses on what players have done in the past; with PCE, we also consider the *future*. In a PCE such as (Cooperate, Cooperate) in Prisoner's Dilemma, players realize that if they deviate from the PCE, then the other player may start to best respond; then, after a while, they will likely end up in some NE, and thus have a payoff that is guaranteed to be no better than (and is often worse than) that of the PCE. Although cooperation here (and in other games) gives a solution concept that is arguably more "fragile" than NE, players may still want to play a PCE because it gives a better payoff. Of course, we are considering one-shot games, not repeated games, so there is no future (or past); however, it is not uncommon that people consider games in a repeated way eventhough it is a one-shot game.

It is easy to see that a PCE does not always exist. Consider the Nash bargaining game [3]. Each of two players suggests a number of cents between 0 and 100. If their total demand is no more than a dollar, then they each get what they asked for; otherwise, they both get nothing. Each pair (x, y) with $x+y = 100$ is a NE, so there is clearly no strategy profile that gives both players a higher payoff than they get in every NE.

Finally, we define PCE for general n-player games. We cannot just use the definition above, since "best response" is not well defined. For example, in a 3-player game, it is not clear what it would mean for players 2 and 3 to make a best response to a strategy of player 1; what might be best for player 2 might not be best for player 3. To deal with this, given an n-player game G and a strategy s_i for player i, let G_{s_i} be the $(n-1)$-player game among the players in other than i that results when player i plays s_i. Let $NE^G(s_i)$ denote the set of Nash equilibria of G_{s_i}. (Note that $NE(s_i) = BR(s_i)$ if there are two players,) We extend the definition of PCE to n-player games for $n > 2$ by replacing $BR^G(s_i)$ by $NE^G(s_i)$ in the definition of $BU^G(s_i)$. PCE is then defined as before, and all its properties are maintained.

3. α-PCE AND M-PCE

α-PCE is a more quantitative version of PCE. Roughly speaking, it takes into account the degree of cooperation exhibited by a strategy profile.

Definition 4. *A strategy profile s is an α-PCE in a game G if $U_i(s) \geq \alpha + BU_i^G$ for all $i \in N$.*

Clearly, if s is an α-PCE, then s is an α'-PCE for $\alpha' \leq \alpha$, and s is a PCE iff s is a 0-PCE. Note that an α-PCE imposes some "fairness" requirements. Each player must get at least α more (where α can be negative) than her best possible outcome if the other players best respond.

Of course, we are interested in α-PCE with the maximum possible value of α.

Definition 5. *The strategy profile s is an* maximum-PCE (M-PCE) *in a game G if s is an α-PCE and for all $\alpha' > \alpha$, there is no α'-PCE in G.*

Every game has an M-PCE; in fact, it has a *Pareto optimal* M-PCE (so that there is no other strategy profile where all players do at least as well and at least one does better). M-PCE does well at predicting behavior in quite a few games of interest. For example, in Prisoner's Dilemma, (Cooperate, Cooperate) is the unique M-PCE; and in the Nash bargaining game, (50, 50) is the unique M-PCE. As the latter example suggests, the notion of an M-PCE embodies a certain sense of fairness. In cases where there are several PCE, M-PCE gives a way of choosing among them.

4. EXAMPLES

The following table compares the payoffs of NE, PCE and M-PCE in a number of games. Besides prisoner's dilemma, the traveler's dilemma, and Nash bargaining, we also consider a coordination game, where players can choose either action a or b, and $U(a, a) = (k_1, k_2)$, $U(a, b) = U(b, a) = (0, 0)$, $U(b, b) = (1, 1)$; we consider other games in the full paper. (In the table below, the M-PCE for the coordination game is the PCE if it exists, and is necessarily one of (1,1) or (k_1, k_2) even if the PCE does not exist, the exact choice depending on k_1 and k_2.)

Problem	NE	PCE	M-PCE
Coordination game	(k_1, k_2) and $(1, 1)$	(k_1, k_2), if $k_1, k_2 > 1$; (1,1), if $k_1, k_2 < 1$; no solution otherwise.	(1,1) or (k_1, k_2)
Prisoner's Dilemma	(1,1)	$(U_1(s), U_2(s))$ for s such that $U_1(s) \geq 1$ and $U_2(s) \geq 1$	(3,3)
Traveller's Dilemma	(2,2)	$(U_1(s), U_2(s))$ for s such that $U_1(s) \geq 99$ and $U_2(s) \geq 99$	(100,100)
Nash bargaining	(0,100), . . . (100,0)	No solution	(50, 50)

Table 1: Comparison of NE, PCE, and M-PCE.

As these examples suggest, M-PCE is quite an attractive solution concept. It always exists, explains cooperative behaviors, and typically provides the high payoffs compatible with what can be viewed as a form of fairness.

5. REFERENCES

[1] K. Basu. The traveler's dilemma: paradoxes of rationality in game theory. *American Economic Review*, 84(2):391–395, 1994.

[2] T. Becker, M. Carter, and J. Naeve. Experts playing the Traveler's Dilemma. Discussion paper 252/2005, Universität Hohenheim, 2005.

[3] J. Nash. The barganing problem. *Econometrica*, 18:155–162, 1950.