

Bargaining efficiency and the repeated prisoners' dilemma

Bhaskar Chakravorti
The Monitor Corporation

John Conley
Vanderbilt University

Abstract

The infinitely repeated prisoners' dilemma has a multiplicity of Pareto-unranked equilibria. This leads to a battle of the sexes problem of coordinating on a single efficient outcome. One natural method of achieving coordination is for the players to bargain over the set of possible equilibrium allocations. If players have different preferences over cooperative bargaining solutions, it is reasonable to imagine that agents randomize over their favorite choices. This paper asks the following question: do the players risk choosing an inefficient outcome by resorting to such randomizations? In general, randomizations over points in a convex set yields interior points. We show, however, that if the candidate solutions are the two most frequently used — the Nash and Kalai–Smorodinsky solutions — then for any prisoners' dilemma, this procedure guarantees coordination of an efficient outcome.

This is a minor correction to a Previously published version.

Citation: Chakravorti, Bhaskar and John Conley, (2004) "Bargaining efficiency and the repeated prisoners' dilemma." *Economics Bulletin*, Vol. 3, No. 3 pp. 1–8

Submitted: February 12, 2004. **Accepted:** February 12, 2004.

URL: <http://www.economicsbulletin.com/2004/volume3/EB-04C70005A.pdf>

1. The Problem

We consider the canonical *prisoners' dilemma* (PD) given by the bimatrix game in Figure 1. Any prisoners' dilemma game (with strictly dominated strategies) can be transformed into a game of the form given below by subtracting the (b, R) payoffs from all cells and normalizing to one the marginal cost of choosing the dominated strategy, given that the opponent chooses the dominant strategy.

The players are 1 (the row player) and 2 (the column player). The payoffs are in terms of von-Neumann Morgenstern utilities and the class of PD games is characterized by $\{(a, d) = (a_1, a_2, d_2, d_2) \in \mathfrak{R}^4 \text{ s.t. } 0 < a_i < d_i \text{ for } i = 1, 2\}$. We shall denote the probability with which player 1 chooses t by p_1 and the probability with which player 2 chooses L by p_2 .

Given (a, d) , an *allocation* is a pair of payoffs $x = (x_1, x_2)$ residing in $\mathcal{X}(a, d)$, the set of all allocations achievable via mixed strategies. With slight abuse of notation, we shall simply write $\mathcal{X}(a, d)$ as \mathcal{X} . The set of all *individually rational allocations* in \mathcal{X} is denoted $IR(\mathcal{X}) \equiv \{x \in \mathcal{X} \mid x \geq (0, 0)\}$. The *Pareto-efficient frontier* of \mathcal{X} is denoted $EFF(\mathcal{X}) \equiv \{x \in \mathcal{X} \mid \forall \hat{x} \in \mathfrak{R}^2, \text{ if } \hat{x} \gg x \text{ then } \hat{x} \notin \mathcal{X}\}$ ¹ It is well known that the unique solution to the PD is $(p_1, p_2) = (0, 0)$.

Consider now the *infinitely repeated PD*, with discount factor $\delta \in [0, 1]$ which we denote (PD^∞, δ) . By the Folk Theorem for infinitely repeated games (for example, see Fudenberg and Maskin 1986) for every x in the interior of $IR(\mathcal{X})$ there exists a $\hat{\delta}$ such that for all $\delta \in (\hat{\delta}, 1]$, there is a subgame perfect equilibrium of (PD^∞, δ) with payoff x in every repetition. Given the multiplicity of equilibria in the game (PD^∞, δ) , it is not immediately clear how agents coordinate their actions to get a specific payoff.

One approach to the problem of multiple equilibrium is to refine the equilibrium

¹ The three vector inequalities are denoted $\gg, >, \text{ and } \geq$.

concept. This is a vast literature. See, for example, Kajii and Morris (1997) or van Damme (2002) for a recent survey of parts of this research. Another alternative is to modify the way that the repeated game is constructed in order to narrow the equilibrium set. See Chakravorti, Conley and Taub for an example of this type of work using probabilistic cheap talk.

In this note, we suppose instead that agents' agree to resolve this coordination problem through the use of bargaining theory. In general, a two person bargaining problem consists of a pair (S, d) , where $S \subset \mathfrak{R}^2$ is interpreted as the set of feasible allocations over the two agents, and $d \in S$, called the disagreement point, is the allocation agents receive if they fail to find a mutually satisfactory way of dividing the surplus. In the case of the repeated prisoners' dilemma game described above, the subgame perfect equilibrium payoffs of stage game can be taken as a set of feasible allocations which converges to the individually rational set in the limit as the discount factor goes to one. For simplicity, we take the limiting case and set $S = IR(\mathcal{X})$. The disagreement point is of course the payoff agents receive when they choose noncooperative strategies. Thus, $d = (0, 0)$ in our case. Note that the $IR(\mathcal{X})$ is closed, convex and comprehensive set which contains elements that strictly dominate the disagreement point. A *bargaining solution* is a function ϕ which maps a class of bargaining problems Σ into \mathfrak{R}_+^2 such that for all $(S, d) \in \Sigma$, it is the case that $\phi(S, d) \in S$.

More formally, we consider a situation in which agents resolve the problem of multiple subgame perfect equilibria using the following procedure.

1. First, a bargaining problem is generated by taking the set of individually rational allocations, $IR(\mathcal{X})$, as the set of feasible payoffs, and the payoff associated with agents jointly playing defection strategies, (b, R) , as the disagreement point.
2. Each player i proposes a bargaining solution, ϕ^i to resolve the bargaining $(IR(\mathcal{X}), 0)$ which the agents face each time the stage game is played. In general,

it is likely that $\phi^i \neq \phi^j$.

3. A public randomization over the alternative solution outcomes $\phi^i(IR(\mathcal{X}), 0)$ and $\phi^j(IR(\mathcal{X}), 0)$ is used to determine the allocation to be implemented. Let $\mathcal{P}(IR(\mathcal{X}), 0)$ denote the allocation thus implemented in any given round of the stage game.

The specific question we are interested in is whether or not there are circumstances under which this procedure will lead to Pareto efficient outcomes.

2. Some Definitions

We begin this investigation by recalling the definitions of the two central solutions in the axiomatic theory of bargaining. The *Nash solution* (Nash, 1950), N is defined as follows:

$$N(S, d) = \underset{\substack{x \in S \\ x \geq d}}{\operatorname{argmax}} (x_1 - d_1)(x_2 - d_2).$$

The *Kalai-Smorodinsky solution* (Kalai and Smorodinsky, 1975), K is found by taking the maximal element in the feasible set on the line connecting the disagreement point and the *ideal point*. Formally, the ideal point is defined as:

$$a(S, d) = \left(\max_{\substack{x \in S \\ x \geq d}} x_1, \max_{\substack{x \in S \\ x \geq d}} x_2 \right),$$

and the Kalai-Smorodinsky solution is defined as:

$$K(S, d) \equiv \{x \in S \mid \exists \lambda \in [0, 1] \text{ s.t. } x = \lambda d + (1 - \lambda)a(S, d) \text{ and } \nexists z \in S, \text{ s.t. } z \gg x\}$$

These two solutions can be characterized by sets of axioms. We now give formal definitions of the axioms which are necessary for our purposes:

A *permutation operator*, π , is a bijection from $\{1, 2, \dots, n\}$ to $\{1, 2, \dots, n\}$. Π^n is the class of all such operators. Let $\pi(x) = (x_{\pi^{-1}(1)}, x_{\pi^{-1}(2)}, \dots, x_{\pi^{-1}(n)})^2$ and $\pi(S) = \{y \in \mathfrak{R}^n \mid y = \pi(x)x \in S\}$.

Symmetry (SYM): If for all permutation operators $\pi \in \Pi^n$, $\pi(S) = S$ and $\pi(d) = d$, then $F^i(S, d) = F^j(S, d) \forall i, j$.

An *affine transformation on \mathfrak{R}^n* is a map, $\lambda : \mathfrak{R}^n \rightarrow \mathfrak{R}^n$, where for some $a \in \mathfrak{R}^n$ and $b \in \mathfrak{R}_{++}^n$, $\lambda(x) = a + bx$. Λ^n is the class of all such transformations. Let $\lambda(S) = \{y \in \mathfrak{R}^n \mid y = \lambda(x), x \in S\}$.

Scale Invariance (S.INV): $\forall \lambda \in \Lambda^n$, $F(\lambda(S), \lambda(d)) = \lambda(F(S, d))$.

Individual Monotonicity (I.MON): For all $(S, d), (S', d') \in \Sigma$, such that $d = d'$ and $S \subseteq S'$ if $a_j(S, d) = a_j(S', d')$ for some $j = 1, 2$, then $\phi_i(S', d') \geq \phi_i(S, d)$.

It is well known that both the Nash and the Kalai-Smorodinsky solutions satisfy the *SYM* and *S.INV* axioms while the Kalai-Smorodinsky solution alone satisfies the *I.MON* axiom.

3. Result

We provide the following answer to the question posed above:

Theorem 1. *If $\phi^i = N$ and $\phi^j = K$, then $\mathcal{P}(\mathcal{X}) \in EFF(\mathcal{X})$.*

Proof/

² Subscripts indicate the components of a vector

Suppose that the theorem were not true. Refer to $(IR(\mathcal{X}), O)$ where $IR(\mathcal{X}) = ABCO$ given in Figure 2. In this case, without loss of generality, suppose that $K(IR(\mathcal{X}), O)$ picks the allocation K and $N(IR(\mathcal{X}), O)$ picks the allocation N on $EFF(\mathcal{X}) \cap IR(\mathcal{X})$. Next, consider the bargaining problem $AECO$ given in Figure 3. Note that E is to the left of C and to the right of K ³. Since $a(AECO, O) = a(IR(\mathcal{X}), O)$, it holds that $K(AECO, O)$ picks the allocation K' . Next, consider the smallest triangle containing $AECO$, (i.e. AFO). By the *I.MON* axiom, $K_1(AFO, O) \geq K_1(AECO, O)$, since $a_1(AFO, O) > a_1(AECO, O)$. In words this means that the Kalai-Smorodinsky outcome of the problem AFO must not lie to the left of K' . By definition, N is the point of contact between two convex sets (i.e. the upper contour set defined by the rectangular hyperbola in \mathcal{X} and the bargaining problem $(IR(\mathcal{X}), O)$ whose interiors are non-intersecting). Since $IR(\mathcal{X}) \subset AFO$, and AFO is convex, we have $N(IR(\mathcal{X}), O) = N(AECO, O) = N(AFO, O)$. By the axioms of *S.INV* and *SY*, given that AFO is triangular, $N(AFO, O) = K(AFO, O)$. Given that N is to the left of K' , we have a contradiction.

■

4. Conclusion

We conclude that for all two person prisoners' dilemma problems, the Nash and Kalai-Smorodinsky solutions suggest allocations on the same facet of the feasible set, and so any randomization between these solutions will also be Pareto efficient.

³ We thank an anonymous referee for pointing out that our previous figure fell outside of the domain of problems we consider.

Thus, while the noncooperative equilibrium concept of subgame perfection leave too many equilibria and no prediction about which of these the agents will eventually settle on, we find that cooperative solution concepts can resolve this problem in a Pareto efficient way even when agents disagree about which solution to choose.

References

- Chakravorti B., J. Conley and B. Taub** (2002): “Probabilistic Cheap Talk,” *Social Choice and Welfare*, vol. 19, pp. 281-294.
- Fudenberg, D and E. Maskin** (1986): “The Folk Theorem in Repeated Games with Discounting and Incomplete Information,” *Econometrica*, vol. 54, pp. 522-54.
- Kajii, A. and S. Morris** (1997): “Refinements and Higher Order Beliefs: A Unified Survey,” Northwestern Center Working Paper No. 1197.
- van Damme, E.** (2002): “Strategic Equilibrium,” in *Handbook of Game Theory with Economic Applications Volume 3*, by R. Aumann and S. Hart, North-Holland, Amsterdam.
- Kalai, E. and M Smorodinsky** (1975): “Other Solutions to Nash’s Bargaining Problem,” *Econometrica*, vol. 43, pp. 513-18.
- Nash, J** (1950): “he Bargaining Problem,” *Econometrica*, vol. 18, pp. 155-62.

PD game		Player 2	
		L	R
Player 1	T	a_1	d_2
	B	d_1	0

Figure 1




Figure 2




Figure 3