

Indefinitely Repeated Games:

A Response to Carroll

Neal C. Becker and Ann E. Cudd

In his article "Indefinite Terminating Points and the Iterated Prisoner's Dilemma", John Carroll asserts that there are "only uncooperative equilibria in finitely, but infinitely, iterated games," and that this "calls into question the significance of the existence of cooperative equilibria in infinitely iterated Prisoner's Dilemma games."ⁱ It is the purpose of this note to show that the claim is overstated for two reasons: (1) the iterated games he refers to are equivalent to finitely and definitely iterated games, of which it is well known that there are only uncooperative equilibria; and (2) his notion of an indefinitely iterated game is without any philosophical interest beyond that of the definitely iterated game.

The problem centers around Carroll's notion of an indefinite iteration, which he characterizes using what he calls a 'terminating \mathbf{p} -function', defined as follows.

\mathbf{p} is a terminating \mathbf{p} -function if and only if

- (1) $\mathbf{p}(t) = 1$, and
- (2) there exists a natural number ω such that
 - (a) $\mathbf{p}(\omega) > 0$, and
 - (b) for all natural numbers $n > \omega$, $\mathbf{p}(n) = 0$.ⁱⁱ

He tells us that we can think of $\mathbf{p}(t)$ as the probability that the game will be iterated t times. Notice that according to this definition the game will not be played more than ω times, so that the sense in which the endgame is indefinite is that it might be played fewer than ω times. Since there is an iteration which will be the last if the game is iterated even that long, it seems to us that it is not at all surprising that a form of the backwards induction proof that shows that there are only noncooperative equilibria in finitely repeated Prisoner's Dilemma (hereafter, PD) games would apply here, as Carroll demonstrates.

We will argue that the indefinite iteration to which Carroll refers is just equivalent to a finitely repeated game with a definite final iteration. This would show what we claim, which is that Carroll's result is a trivial extension of an already well-known result. The argument goes as follows. Suppose we have the prisoner's dilemma of the following matrix:

		player 2	
		C	D
player 1	C	x, x	z, y
	D	y, z	w, w

Figure 1

The payoffs are real numbers related as follows: $y > x > w > z$. Carroll defines a function $\mathbf{p}^*(t)$ to represent the probability that the game is iterated in period $t+1$.ⁱⁱⁱ Let $\langle a, b \rangle X_{i,t}$ be the payoffs to player i in period t given that i plays supergame strategy a and the other player plays supergame strategy b . The iterated supergame payoff is then given by the sum of the payoffs in each game given the

strategy pairs of the players multiplied by the probability that they will be played:

$$p^*(t-1) \langle a, b \rangle X_{i,t}$$

Note that this is just equivalent to:

$$(*) \quad p^*(t-1) \langle a, b \rangle X_{i,t} ,$$

since ω is the last period with a positive probability of being played. The '[infinity sign]' is gratuitous here, and serves only to cloud the issue of whether the game is to be repeated finitely but indefinitely (i.e. with a finite upper bound on the number of possible iterations) or infinitely but indefinitely (that is, with no finite upper bound on the number of possible iterations). Note also that the utility functions which Carroll is presupposing must be von Neumann-Morgenstern utility functions, since he assumes in writing the iterated game payoff functions this way that they have the expected utility property.^{iv}

Finally note that Carroll seems to be supposing that the players have common knowledge of the iterated supergame, in particular, of ω , and of the players' rationality.^v

The problem, then, is to find the strategies a and b such that equation (*) is maximized for each player, given the payoff matrix. By the expected utility property, we can write the utility of a stochastic game as the utility for the same game played with certainty but with the payoffs multiplied by the probability of the game being played. So the problem is equivalent to finding the strategies a and b which maximize the equation:

$$\langle a, b \rangle X^*_{i,t}$$

where X^* represents the payoffs of the following matrices for each period:

		player 2		
		C	D	
player 1	C	$\mathbf{p}^*(0) (x, x)$	$\mathbf{p}^*(0) (z, y)$	Period 1
	D	$\mathbf{p}^*(0) (y, z)$	$\mathbf{p}^*(0) (w, w)$	

		player 2		
		C	D	
player 1	C	$\mathbf{p}^*(1) (x, x)$	$\mathbf{p}^*(1) (z, y)$	Period 2
	D	$\mathbf{p}^*(1) (y, z)$	$\mathbf{p}^*(1) (w, w)$	

.

.

.

		player 2		
		C	D	
player 1	C	$\mathbf{p}^*(\omega-1) (x, x)$	$\mathbf{p}^*(\omega-1) (z, y)$	Period ω
	D	$\mathbf{p}^*(\omega-1) (y, z)$	$\mathbf{p}^*(\omega-1) (w, w)$	

Figure 2

Furthermore, since these are von Neumann-Morgenstern utilities, we can make a linear transformation to each matrix, which can be chosen for each matrix so that this supergame is equivalent to ω iterations of the original matrix, without changing the iterated game equilibrium strategy. (Just multiply the payoffs of

the k th matrix by $1/p^*(k-1)$, which must exist, since by definition $p^*(k) > 0$ for all $1 \leq k < \omega$.) Thus Carroll's indefinitely iterated PD supergame is equivalent to the definitely finitely iterated PD supergame, and so it is no surprise to learn that it has only an uncooperative solution.

It is still an interesting question, however, whether situations of indefinitely repeated interaction should be modeled as finitely repeated games or as the standard sort of indefinitely repeated games. We would like to conclude by giving a few reasons to think that the indefinitely repeated variety is better suited to the task of modeling a risky future. The 'standard indefinitely iterated game', as we shall call it, is the game represented by Figure 1 repeated with probability d each time, so that the probability that the game will be repeated once (i.e. played twice) is d , ($0 < d \leq 1$), the probability that it will be repeated twice is d^2 , and in general, the probability that iteration t of the game, g_t , will be played is d^{t-1} . The difference between this and Carroll's finite but indefinitely iterated game is that the standard iterated game has a small but positive probability of continuing in any finite iteration. But since it is indeed certain that we are all mortal, and thus face only finitely many iterations of any decision situation, Carroll claims that the iterated game with a definite upper bound on the number of iterations is a better model of reality. He writes, "my definition of the iterated payoff is specifically designed to capture the finiteness of genuine iterated Prisoner's Dilemmas."^{vi}

There are two problems with Carroll's claim. First, as we have shown, the definiteness induced by the upper bound on the number of iterations in Carroll's game, (i.e. clause (ii) of his definition of the terminating p -function), overwhelms any indefiniteness in the model, and makes it equivalent to a finitely

and definitely iterated game. That is, by introducing a definite endpoint Carroll has trivialized the indefinite nature of the iterations. Second, in defining his terminating p -function as a finite sequence of probabilities which sum to 1, (clause (i) of his definition), he has effectively fixed the conditional probability of future iterations of the game beyond ω . In particular, the conditional probability of playing another game given that one is playing the ω th iteration is 0. But is it realistic to suppose that one is usually, or perhaps ever, certain that the current interaction is the last one? We think that it is more realistic to imagine that in any interaction there is a positive probability that one will experience another interaction; that's why it is a bad idea (unless it is a rationally chosen strategy to precommit yourself) to burn your bridges.

An analogy may make our point clearer. Suppose Ethel lives longer than anyone has ever lived before, say 175 years. We might imagine that a game theorist modeling PD iterations in life would set the probability of living 176 years at 0. But is it reasonable for Ethel to believe that it is certain that she will not live another year? Given that she has lived 175 years, it seems to us that she would be justified in placing a positive, if small, probability on living another year. The same point could be made, we believe, for any interaction situation -- it seems to us that it is normally unreasonable to believe with certainty that any particular interaction is the last of its kind. Carroll claims that it is certain that we will not face a decision situation for millions of centuries, but in order to build this certainty in he must set an arbitrary upper bound on the number of iterations, and this is problematic for at least the last game. The standard indefinitely iterated game, on the other hand, would place

a very small, vanishing to zero, probability on the chances of interactions continuing for "millions of centuries".

An alternative approach, which captures the idea that the probability of continuing iterations of the game decreases with time, but doesn't require the conditional probability of the game continuing at any stage to be zero, is to let the probabilities decrease monotonically with time reaching zero only in the limit, if at all. Let $p'(t)$ be the probability that the game will be played t times, such that p' has two properties:

(1') $\lim_{t \rightarrow \infty} p'(t) = [\alpha]$, $1 \geq [\alpha] \geq 0$, $p'(t)$ monotonically decreasing;

(2') $1 \geq p'(t) > 0$.

Notice that for $p'(t)$ a constant 0 we have the one-shot PD, for $p'(t)$ a constant 1 we have the infinitely repeated PD, and for $p'(t)$ we have the standard indefinitely repeated PD. From the results of the standard indefinitely iterated PD we know that for $[\alpha]$ small enough (relative to w, x, y, z) there exists a finite iteration t^* such that the probability of continuing is very close to $[\alpha]$ and it no longer pays the players to cooperate, and this is enough to begin the backwards induction to show that they will never cooperate for such an $[\alpha]$. Thus for some such games Carroll's result would hold despite the potentially infinite iteration.

Of course, there may be artificial situations which are best modelled by the indefinitely finitely iterated PD, such as an Axelrod-type tournament in which an upper bound on the number of iterations has been set. But these situations must be manufactured to have the definite upper bound; they are not naturally occurring. And again, such games are equivalent to finitely repeated games.

Notes

i. John Carroll, (1987), pp.255-6.

ii. Ibid., p.249.

iii. Carroll claims on p.251 that " $p^*(t)$ could be interpreted as the probability that game [t] will be played," but we take it that it should read ' $p^*(t-1)$ ', given his definition of it in terms of $p(t)$.

iv. See John Harsanyi, (1977), p.33, for a discussion of utility functions and the expected utility property.

v. See David Kreps, et.al., (1982), pp.245-252. They show that there can be cooperative equilibria when there is not common knowledge of the players' options or motivation.

vi. Carroll, op. cit., p.250.

References

Carroll, J.W.: 1987, 'Indefinite Terminating Points and the Iterated Prisoner's Dilemma', Theory and Decision, **22**, 247-256.

Harsanyi, J.C: 1977, Rational Behavior and Bargaining Equilibrium in Games and Social Situations, Cambridge University Press, Cambridge.

Kreps, D.M., Milgrom, P., Roberts, J., and Wilson, R.: 1982, 'Rational Cooperation in the Finitely Repeated Prisoner's Dilemma', Journal of Economic Theory, **27**, 245-252.

Neal C. Becker
Department of Economics
University of Kansas
Lawrence, Ks. 66045
(913) 864-3501

Ann E. Cudd
Department of Philosophy
University of Kansas
Lawrence, Ks. 66045
(913) 864-3976