Acta Psychologica 187 (2018) 1–8

Contents lists available at ScienceDirect

Acta Psychologica

journal homepage: www.elsevier.com/locate/actpsy

Persistent cooperation and gender differences in repeated Prisoner's Dilemma games: Some things never change

Andrew M. Colman*, Briony D. Pulford, Eva M. Krockow

Department of Neuroscience, Psychology and Behaviour, University of Leicester, Leicester LE1 7RH, UK

ARTICLE INFO

PsycINFO classifications:
2340
2970
3020*Keywords:*
Cooperation
Endgame effect
Gender difference
Prisoner's Dilemma
Social dilemma

ABSTRACT

In the finite-horizon repeated Prisoner's Dilemma, a compelling backward induction argument shows that rational players will defect in every round, following the uniquely optimal Nash equilibrium path. It is frequently asserted that cooperation gradually declines when a Prisoner's Dilemma is repeated multiple times by the same players, but the evidence for this is unconvincing, and a classic experiment by Rapoport and Chammah in the 1960s reported that cooperation eventually recovers if the game is repeated hundreds of times. They also reported that men paired with men cooperate almost twice as frequently as women paired with women. Our conceptual replication with Prisoner's Dilemmas repeated over 300 rounds with no breaks, using more advanced, computerized methodology, revealed no decline in cooperation, apart from endgame effects in the last few rounds, and replicated the substantial gender difference, confirming, in the UK, a puzzling finding first reported in the US in the 1960s.

1. Introduction

The archetypal social dilemma is the two-player Prisoner's Dilemma (PD), a game that has been subjected to much experimental investigation in the history of experimental games and behavioral game theory (Rapoport, Seale, & Colman, 2015; Roth, 1995). Among the many reasons for its enduring popularity is the fact that it provides a conceptual structure within which phenomena such as cooperation and competition, trust and trustworthiness, altruism and spite, threats, promises, commitments, and collective rationality can be formalized and investigated rigorously, on the basis of behavioral measures rather than mere questionnaire responses (Pruitt & Kimmel, 1977; Rapoport & Chammah, 1965a); but what attracts researchers to it more than anything else is its paradoxical character and the challenge of explaining why players cooperate. In single-play (one-shot) and finite-horizon repeated PD, there are compelling arguments, explained in the Section 1.1 below, why rational players should never cooperate.

The research reported in this article focuses on cooperation in the finite-horizon repeated PD, in which the game is repeated over a finite number of rounds by the same players, who know in advance how many rounds will be played. Many experiments have addressed this issue, but the vast majority used only short sequences of repetitions. An exhaustive meta-analysis of experiments on trust and cooperation in both two-player and multi-player social dilemmas (Balliet & Van Lange, 2013) found 212 experiments, 132 using one-shot interactions and the rest only small numbers of repetitions ($M = 6.07$, $SD = 13.54$); hardly any used more than 50 repetitions (see also Balliet, Mulder, & Van Lange, 2011; Embrey, Fréchette, & Yuksel, 2018).

The most ambitious experiment with long sequences or repetitions (Rapoport & Chammah, 1965a) involved 140 experimental subjects playing 300 rounds, and seven different PDs. The researchers reported an initial decline in cooperation followed by a recovery after many repetitions: "The most typical feature of the time course of a Prisoner's Dilemma protocol is the initial decline in cooperation, followed eventually [after 30–60 rounds] by a recovery" (p. 200). However, there are some aspects of the experiment that make this conclusion difficult to interpret. The instructions given to the subjects began: "You will be playing a game," probably priming an initially competitive mental set, because the objective in virtually all familiar indoor and outdoor games is to beat the opponent. The experiment included incentive payments, but the level of remuneration was 1/10 of a penny (US cent) per payoff point, derisory even in the 1960s. Above all, the "time courses"—the claimed declines and recoveries in the relative frequency of cooperative choices—were averaged over seven PDs with different payoffs and presented only graphically as moving averages. In Figure 7 (p. 90) and Figure 17 (p. 97), showing results for the relevant "pure matrix" treatment conditions, it is far from obvious that the reported initial decline is statistically significant, and no evidence is provided to back this up, because appropriate statistical techniques for analyzing time series had not yet been developed when the experiment was conducted.

* Corresponding author.
E-mail addresses: amc@le.ac.uk (A.M. Colman), bdp5@le.ac.uk (B.D. Pulford), emk12@le.ac.uk (E.M. Krockow).

https://doi.org/10.1016/j.actpsy.2018.04.014
Received 20 December 2017; Received in revised form 15 April 2018; Accepted 23 April 2018

(a)

I \ II	C	D
C	3, 3	0, 5
D	5, 0	1, 1

(b)

I \ II	C	D
C	4, 4	0, 5
D	5, 0	1, 1

(c)

I \ II	C	D
C	0, 0	-2, 1
D	1,-2	-1,-1

(d)

I \ II	C	D
C	*R, R*	*S, T*
D	*T, S*	*P, P*

Fig. 1. Prisoner's Dilemma games with different payoff and index of cooperation (K) values. (a) Conventional version with $K = 2/5$ or 0.40; (b) Mild version with a higher index of cooperation, $K = 3/5$ or 0.60; (c) Original symmetric version from Tucker (1950/2001) with $K = 1/3$ or 0.33 approximately; (d) Generalized payoff matrix for any symmetric 2×2 game. The index of cooperation $K = (R - P)/(T - S)$.

The same classic study was also the first to report a large and unexpected gender difference, with male/male pairs cooperating almost twice as frequently as female/female pairs (Rapoport & Chammah, 1965b). Such large gender differences are seldom reported in psychology, and a natural expectation, based on traditional sex roles and socialization, would be of more cooperation in female/female than male/male pairs. The experiment reported below was designed as a conceptual replication to check Rapoport and Chammah's findings using more rigorous experimental and data-analytic techniques and also to provide some evidence on the cross-cultural generalizability and temporal stability of these findings.

1.1. Theoretical considerations

Fig. 1(a) shows the payoff matrix of a PD with payoff values originally introduced by Scodel, Minas, Ratoosh, and Lipetz (1959), popularized by Axelrod (1980a, 1980b, 1984), and nowadays frequently described as "conventional" (Press & Dyson, 2012). The original symmetric version used by Tucker (1950/2001) when he named the newly discovered game in 1950 is shown in Fig. 1(c). Player I chooses between the rows marked C (cooperate) and D (defect), Player II independently chooses between columns C and D, and the cell in which the pair of strategy choices intersect is the outcome of the game, with the payoffs to Player I and Player II listed in that order by convention.

In the conventional version (Fig. 1a), if the game is played just once, then both players do better if both cooperate (each receiving 3 units) than if both defect (each receiving 1 unit). Nevertheless, rational players are bound to defect, because D is a *dominant strategy* for both players, yielding a higher payoff than the C strategy whether the co-player chooses C or D, and D is therefore an unconditionally best strategy. The (D, D) outcome, in which both players choose their optimal D strategies, is the unique *Nash equilibrium* of this game—the only outcome in which each player's strategy is a best reply to the co-player's, in the sense that no other strategy yields as high a payoff against the co-player's chosen strategy. For example, the (C, D) outcome is out of equilibrium: Player I's choice of C is not a best reply to Player II's D, because Player I could have received a better payoff by choosing D, given Player II's choice of D. Only in the (D, D) outcome are both players strategies best replies and hence in Nash equilibrium.

In a repeated PD with no finite horizon or end-point known in advance, there are reasons to cooperate in spite of the dominance of the D strategy in the one-shot version, because rounds that have yet to be played cast a "shadow of the future" over earlier rounds. If Player I defects in Round t, then Player II may retaliate with defection in Round $t + 1$ or later, reducing Player I's payoff. But in a repeated PD with a finite horizon—one in which a finite number of rounds are to be played and the players know this number—rational players will defect in every round. This is persuasively proved by the following argument (Luce & Raiffa, 1957, pp. 97–102; Sobel, 1993). Suppose the players know that there are to be exactly 100 rounds. In Round 100, there is no reason to cooperate, because there are no rounds to follow and therefore no possibility of retaliation; therefore, both players will defect in Round 100, because the D strategy is dominant, and both therefore do better by choosing D than C irrespective of what the co-player chooses. In Round 99, both players know that the outcome of Round 100 is predetermined, for the reason just given, therefore there is no reason to cooperate in Round 99, and players will choose their dominant strategies. This argument unfolds backwards in the same way, mandating defection in every round, including the first. Joint defection in every round is the only Nash equilibrium of the finite-horizon repeated PD, and it is proved by the argument above, called *backward induction*. However, the conclusion relies on full common knowledge of rationality. In a highly cited article, Kreps, Milgrom, Roberts, and Wilson (1982) showed that if both players are strictly rational payoff maximizers, but at least one believes that there is even a tiny probability that the other is irrational, then rational cooperation can occur until close to the final round (for slightly different approaches, see Ambrus & Pathak, 2011; Dijkstra & Van Assen, 2017).

1.2. Experimental evidence

The first experimental study of the finite-horizon repeated PD was performed by Dresher and Flood in January 1950 and reported in a RAND research memorandum RM-789 in 1952, subsequently revised and condensed for publication by Flood (1958). Two research subjects, who were friends and had a knowledge of game theory, played exactly 100 incentivized rounds of an asymmetric PD in which $T > R > P > S$ for each player considered separately (see Fig. 1d). The relative frequency of C choices was 73%, and "there was a decided tendency to start with [(D, D)] and then to shift to [(C, C)] rather consistently after about thirty trials" (pp. 14–15), except for the very last round, in which both players defected.

In an influential monograph on game theory, Luce and Raiffa (1957) suggested that repeated PDs should evolve toward joint cooperation in that way: "We feel that in most cases an unarticulated collusion between the players will develop. … This arises from the knowledge that the situation will be repeated and that reprisals are possible" (p. 101). This prediction appeared to be comprehensively refuted when the first full-scale, incentivized experiment was published two years later (Scodel et al., 1959). Introducing for the first time the conventional version of the game shown in Fig. 1(a), their 22 player pairs (all men) completed 50 rounds of the game, and only two pairs showed evidence of collusion or increase in joint cooperation. Overall, significantly more D choices and joint defection outcomes (DD *lock-ins*) were observed in the last 25 rounds than in the first. Minas, Scodel, Marlowe, and Rawson (1960) and others replicated this effect, and an early review of published PD experiments concluded: "In general, the percentage of cooperative responses … tends to decrease over a series of trials" (Gallo & McClintock, 1965, p. 74). This finding has been replicated in more recent research; for example, Cooper, DeJong, Forsythe, and Ross (1996) reported: "Cooperation rates are positive and generally declining over time in the [finite-horizon repeated PD]" (p. 200). A simple learning model predicts just such a decline (Bornstein, Erev, & Goren, 1994), but the experimental studies that have shown a decline have not used long sequences of repetitions (e.g., Cooper et al. used only 10), and the declines reported in empirical studies may have been mere endgame effects as cooperation tends to disappear in the last few rounds.