When Cooperation Extends the Shadow of the Future*

Ivan Dedyukhin†

October 19, 2026

Abstract

I introduce an indefinitely repeated prisoner's dilemma in which the continuation probability depends on how many players cooperate. Theoretically, the payoff from mutual cooperation and the continuation probability it induces act as substitutes in sustaining cooperative equilibria. I test this in a laboratory experiment with a 2×2 between-subjects design that independently varies these two parameters. Three results emerge. First, an increase in the continuation probability associated with joint cooperation promotes cooperation only when it makes cooperation sustainable in equilibrium. Second, early in the experiment, payoff incentives raise cooperation more than theoretically equivalent continuation incentives; this asymmetry disappears with experience. Third, estimated strategies display substitutability between the two forces once treatments are matched on strategic uncertainty.

JEL Classification: C72, C73, C92, C91;
Keywords: repeated game, strategy, prisoner's dilemma, cooperation, endogenous continuation, endogenous discounting, strategic uncertainty, laboratory experiment.

*I am grateful to Ala Avoyan and Daniela Puzzello for their continuous support and guidance throughout this project. I thank Eungik Lee, Natalie Lee, Guillaume Frechette, Jimmy Walker, Yaroslav Rosokha, Tim Cason for helpful comments and suggestions. I also thank participants of the NEXT workshop, internal seminars at Indiana University, the Ninth Biennial Meeting of the Social Dilemmas Working Group, the Economic Science Association Meetings, the Midwest Experimental Economics Meetings for Graduate Students, and Caltech Summer Program in Theory-Based Experiments for valuable feedback. This research was approved by the Indiana University Institutional Review Board (IRB #25757) and pre-registered on AsPredicted (#219709). AI tools (including large language models) were used to assist with language editing and code development. All errors are my own.

†Indiana University (email: idedyukh@iu.edu)

1 Introduction

In standard indefinitely repeated games, current actions affect current payoffs but not the probability of future interaction. Yet in many repeated relationships, cooperation can itself make future interaction more likely: successful cooperation may preserve a relationship, while a breakdown in cooperation may shorten or terminate it. In such settings, cooperation affects not only what players earn today, but also whether they have opportunities to interact tomorrow. During the Cold War, each round of arms negotiations resembled a prisoner's dilemma: unilateral restraint exposed a country to exploitation, while mutual restraint generated immediate gains. But mutual restraint lowered the risk of a catastrophic conflict that would have ended the relationship altogether. The same feedback arises in cartels, where price discipline builds resilience against the demand shocks; in R&D alliances, where sustained joint investment strengthen robustness against the funding and market disruptions that would otherwise leave the partnership unable to continue; and in environmental agreements, where sustained compliance enhances lower probability of environmental shocks.[1] In such settings the shadow of the future is not a fixed feature of the environment: cooperation itself lengthens it. Yet research on repeated games has focused almost exclusively on environments in which the expected duration of the interaction is independent of play. This paper studies cooperation when the continuation of the relationship is itself endogenous to players' actions.

I introduce a variant of the indefinitely repeated prisoner's dilemma (IRPD) in which the probability of continuation depends on current cooperation. After each round, the game continues with probability $\delta(c)$, where $c \in \{0, 1, 2\}$ is the number of players who cooperated in that round. Cooperation today thus generates a "longer tomorrow": it raises the expected duration of the relationship and, with it, the value of all future cooperative gains. Two theoretical observations organize the paper. First, the stage-game reward from mutual cooperation and the continuation probability following mutual cooperation act as substitutes in sustaining cooperation. Cooperation is supportable as a subgame perfect equilibrium whenever $\delta(2)$ exceeds a threshold that is decreasing in the reward, so a higher continuation probability can compensate for a lower reward and vice versa. The economics behind the substitution is that the continuation probability rule itself punishes defection: a defection shortens the expected horizon before the opponent responds, so less of the discipline has to come from larger rewards. Second, endogenous continuation reshapes incentives in a way that exogenous discounting cannot. For some parameter values of $\delta(0)$ and $\delta(1)$ cooperating unilaterally against a defector prolongs the expected life of the relationship. Hence, this new dynamic can make permanent joint defection unsustainable in equilibrium even though defection remains the dominant action in every stage game. For the same reason, because unilateral defection lowers the continuation probability from $\delta(2)$ to $\delta(1)$, unconditional cooperation can be sustained in equilibrium for some parameter values, even though cooperation is strictly dominated in the stage game. Cooperation is

[1] Another example is a public media organization whose repeated fundraising success builds financial resilience against adverse shocks, lowering the probability of closure.

an investment: it gives up the gain from defecting today for a higher probability that the relationship continues.

The empirical part of the paper is organized around two questions. First, do subjects internalize cooperation-contingent continuation? In other words, does cooperation increase when mutual cooperation extends the expected duration of the match, holding the stage game fixed? Second, how does the behavioral effect of continuation incentives compare with that of payoff incentives of equal theoretical value: when the expected payoff from mutual cooperation is held constant, do the two channels generate similar cooperation rates, or is one more effective than the other?

To address these questions, I implement a laboratory experiment in which I independently vary two key parameters: (i) the reward from mutual cooperation and (ii) the continuation probability following mutual cooperation. This yields a 2×2 between-subjects design. A key advantage of this design is that it isolates variation in the expected payoff from cooperative strategies,[2] while holding all other aspects of the game environment constant. As a result, the design allows for a direct comparison of the behavioral effects of increasing the reward from mutual cooperation and increasing the likelihood of future interaction following mutual cooperation.

Three main results emerge. First, subjects respond to cooperation-contingent continuation probabilities: cooperation increases when a higher continuation probability after joint cooperation makes cooperative play supportable as an equilibrium, but not when cooperation is already sustained. This pattern suggests that subjects internalize the equilibrium implications of the continuation structure rather than fully exploiting the additional increase in the expected duration of interaction. Second, early in the experiment, increasing the reward from mutual cooperation raises cooperation, whereas increasing the continuation probability has no effect. With experience, this asymmetry fades: average behavior converges toward indifference, while substantial heterogeneity across sessions remains, with patterns consistent with the basin of attraction of Always Defect. Third, strategy estimation reveals that the two treatments generating the same level of strategic uncertainty produce indistinguishable strategy distributions, so that I cannot reject substitution of effects.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 introduces the game and outlines the theoretical framework. Section 4 describes the experimental design and procedures. Section 5 presents the experimental results. Section 6 concludes.

2 Literature

The present study contributes to the experimental literature on IRPD games, which investigates the extent to which subjects respond to dynamic incentives (e.g., Aoyagi et al., 2019, 2024; Bigoni et al., 2015; Blonski et al., 2011; Dal Bó, 2005; Dal Bó and Fréchette, 2011; Fudenberg et al., 2012; Gill and Rosokha, 2024; Lugovskyy et al., 2020; Proto et al., 2019).[3] This body of work has robustly

[2] That is, strategies that begin with cooperation, such as Grim Trigger or Tit-for-Tat.

[3] For a comprehensive survey of this literature, see Dal Bó and Fréchette (2018).

documented that, on average, subjects adjust their behavior in response to variations in continuation probability and stage-game payoffs in a manner broadly consistent with theoretical predictions. Specifically, cooperation becomes more prevalent when the continuation probability increases or when the reward for mutual cooperation is higher. In prior studies, continuation probabilities were exogenously fixed across histories of play, so higher continuation probabilities typically increased the expected continuation value for all strategies. By contrast, in the present design, the increase in the continuation probability is contingent on mutual cooperation, thereby affecting only the continuation value following successful cooperation. As such, this mechanism acts as a substitute for increasing the reward from cooperation in the stage game.

To the best of my knowledge, the experiment most closely related to the present study is Vespa and Wilson (2019), who investigate action-dependent transitions in a Dynamic Prisoner's Dilemma.[4] In their design, mutual cooperation or defection is required for transition to a new game with a different payoff structure.[5] While their study examines action-dependent transitions between different payoff states, the present design holds the stage game fixed and varies only the probability that the relationship continues. A related design, Kloosterman (2020), embeds an IRPD in a stochastic state that determines whether cooperation is supportable in equilibrium; because the state evolves independently of players' actions, cooperation there cannot itself raise the probability that equilibrium play remains supportable, as it does in the present design.

A related literature studies environments in which continuation is endogenous through a player's choice to end the relationship, rather than through a rule tied to realized play. Ghosh and Ray (1996) and Fujiwara-Greve and Okuno-Fujiwara (2009) model repeated Prisoner's Dilemmas in which players can unilaterally separate and rematch with a new partner, so that cooperation is what makes a partner worth keeping.[6] Experimentally, Wilson and Wu (2017) and Lee (2020) show that allowing players to terminate relationships affects cooperation.[7] In these environments, continuation is determined by players' choices. In the present design, by contrast, continuation is a mechanical function of the realized action profile.

Moreover, this paper contributes to the literature on infinitely repeated games with endogenous discounting. The traditional experimental implementation of discounted repeated games employs a random termination procedure, first proposed by Roth and Murnighan (1978), which matches the expected number of repetitions to the discount factor. Although infinitely repeated games with payoff discounting are theoretically isomorphic to randomly terminated games under the assumption

[4]The Dynamic Prisoner's Dilemma (DPD) extends the standard infinitely repeated PD game by adding a single additional state, so that the stage-game can change across periods.

[5]See also Vespa (2020) for a repeated common-pool resource game in which the endogenously determined stock in one period affects the stage-game payoffs in the next, and Battaglini et al. (2016) for a repeated public good game with accumulation of durable public goods.

[6]See also Rob and Yang (2010) and Fujiwara-Greve et al. (2012).

[7]Relatedly, Kamei (2019) studies endogenous matching institutions in an indefinitely repeated public-goods game, allowing subjects to choose between partner and random matching.

of risk neutrality, they represent distinct environments in practice.[8] Related work shows that how the horizon is implemented also matters for behavior: Mengel et al. (2022) find that the realized length of a match, not just its expected length, affects cooperation.[9] The idea that conduct shapes the likelihood of future interaction dates to Thomas and Feldman (1988), who let the continuation probability depend on whether a single actor, the perceived "leader," cooperates or defects. The present game instead ties continuation to the realized action profile, which depends on both players. A more recent theoretical predecessor is Kochov and Song (2023), who develop a model in which a player's discount factor depends on the stream of payoffs received, and characterize how increasing and decreasing marginal impatience (IMI and DMI) reshape the Pareto frontier and the incentives to sustain cooperation.[10] Their analysis is purely theoretical and does not provide experimental evidence on whether subjects respond to endogenous discounting. Moreover, as I show in Appendix D, the present game falls outside both the IMI and DMI frameworks. The present paper therefore contributes along two dimensions relative to Kochov and Song (2023): it introduces a distinct form of endogenous continuation – one in which cooperation directly raises the probability of future interaction – and it provides the first experimental test of behavior under cooperation-dependent continuation in the infinitely repeated Prisoner's Dilemma.

3 Game

An indefinitely repeated prisoner's dilemma game is a dynamic game that allows us to study how the prospect of future interaction can sustain cooperation (see Abreu et al., 1990; Axelrod, 1984; Friedman, 1971; Fudenberg and Maskin, 1986).

Let $N = \{1, 2\}$ be the set of players. For each player $i \in N$, the action set is $A = \{C(\text{ooperate}), D(\text{efect})\}$, and the action profile at any stage is $a = (a_1, a_2)$, where $a_i \in A$ for each player. The payoff for player i in the stage game is given by $u_i : A^2 \to \mathbb{R}^+$, where $u_i(a)$ denotes player i's payoff from action profile a. The stage-game payoffs $u_i(a^t)$ are symmetric and can be represented in a game matrix in Table 1, where the first (second) entry in each cell corresponds to the payoff of the row (column) player.

Table 1: Prisoner's dilemma

	C	D
C	R, R	S, T
D	T, S	P, P

[8] Indeed, depending on the context and modeling assumptions, either or both interpretations have been adopted; see Mailath and Samuelson (2006, Sect. 4).

[9] See also Lugovskyy et al. (2017) and Anderson et al. (2024), who compare random and finite stopping rules in repeated public goods games.

[10] See also Sekiguchi and Wakai (2016) with a duopoly model with gain/loss asymmetric discounting.

For this game to be a Prisoner's Dilemma, it must be that $T > R > P > S$, and also I require $2R > T+S$, so the alternation between cooperation and defection is not more profitable than mutual cooperation. With such payoff, stage game dominant action is defection, while it is Pareto dominated by mutual cooperation.

The stage game is repeated for an indefinite number of rounds (following standard terminology, I refer to each such repeated interaction as a supergame), and in every round $t = 0, 1, 2, \dots$ the game continues for an additional round with probability $\delta : \{0, 1, 2\} \to [0, 1)$, where $\delta(c)$ denotes probability of next round when there are c players who cooperated. The number of players who cooperated is given by $c : A^2 \to \{0, 1, 2\}$, where $c(a)$ denotes the number of players who cooperated in action profile a:

$$c(a) = \#\{a_i = C \text{ for } i \in \{1, 2\}\}.$$

So, the continuation probabilities are defined as follows: $\delta(0)$ when both defect, $\delta(1)$ when exactly one cooperates, and $\delta(2)$ when both cooperate. Continuation probability $\delta(c)$ is an increasing function, so the relationship between these probabilities is given by:

$$\delta(0) \leq \delta(1) \leq \delta(2).$$

In each period t, players simultaneously choose actions from their respective action sets. The outcome of each period is determined by the action profile chosen.

Assuming risk neutrality, the total expected payoff v_i for player i in the indefinitely repeated game is the expected sum of the payoffs across all periods:

$$v_i(a^0, a^1, \dots) = u(a^0) + \sum_{T=1}^{\infty} \left(\prod_{t=1}^{T} \delta(c^{t-1}) \right) u(a^T),$$

where a^t is the action profile chosen in period t and c^t is a corresponding $c(a^t)$. Also it can be represented recursively:

$$v_i(a^0, a^1, \dots) = u(a^0) + \delta(c^0) v_i(a^1, a^2, \dots).$$

A history h^t up to period t is the sequence of past action profiles: $h^t = (a^0, a^1, \dots, a^{t-1})$. A strategy for player i is a function $s_i : H^t \to A$ that specifies the action a_i^t in period t based on the history h^t, where H^t denotes the set of all possible histories up to period t.

Some of the key supergame strategies in our analysis are as follows.

1. Always Cooperate (AC) and Always Defect (AD) are the strategies that choose C and D, respectively, for every history.
2. s_i is *Grim* if $a_i^t(h^t) = C$ if $h^t = ((C, C), \dots, (C, C))$ or $\varnothing$, and $a_i^t(h^t) = D$ otherwise.
3. s_i is TFT if $a_i^t(h^t) = C$ if $a_j^{t-1} \in \{C; \varnothing\}$ and $a_i^t(h^t) = D$ otherwise.

AD can be supported as subgame perfect Nash equilibrium (SPE) under the following condition:

$$\frac{1-\delta(1)}{1-\delta(0)} \cdot P \geq S. \tag{AD}$$

Hence, defection may not be sustained as a SPE for certain parameter combinations. The intuition is as follows: a unilateral deviation to cooperation increases the continuation probability (from $\delta(0)$ to $\delta(1)$), thereby raising the expected value of future payoffs. If this increase in the expected continuation value outweighs the cost of receiving the sucker's payoff S instead of the punishment payoff P, then defection is no longer an optimal strategy. In such cases, players may find it preferable to deviate from unconditional defection, undermining the sustainability of the Always Defect strategy as an equilibrium.

Cooperation can be supported as a subgame perfect equilibrium (SPE) under the *Grim* strategy if condition AD holds and the following additional condition is satisfied:

$$\frac{R}{1-\delta(2)} \geq T + \delta(1)\frac{P}{1-\delta(0)} \tag{Grim}$$

Here the stage-game reward from joint cooperation and the continuation probability following mutual cooperation act as substitutes in sustaining cooperation. Both R and $\delta(2)$ exist in the left-hand side of the equation Grim, so the right-hand side of the equation remain unchanged, so that by increasing R and decreasing $\delta(2)$ we can keep the equation exactly the same, keeping the incentive of the *Grim* strategy and expected cooperative payoff. Figure 1 shows how the condition on $\delta(2)$ for sustaining cooperation is decreasing with the R, so a higher R can be compensated with the lower $\delta(2)$. I explore those incentives in the following experiment.

Additionally, this structure creates a theoretical environment in which cooperation can be supported in games with a shorter expected length, while failing to be an equilibrium in games with a longer expected horizon. Consider, for example, a Prisoner's Dilemma with payoffs $T = 50$, $R = 32$, $P = 25$, and $S = 12$. In the standard indefinitely repeated setting with a fixed continuation probability, the *Grim* strategy constitutes a SPE if the continuation probability exceeds 0.72. Therefore, if the continuation probability is set to 0.65, cooperation cannot be sustained as a SPE. However, in the endogenous continuation game introduced above, using the same stage-game payoffs, cooperation can be supported as a SPE under a continuation rule such that $\delta(2) = 0.6 > \delta(1) = 0.3 > \delta(0) = 0.2$. In this case, the expected length of the game is strictly shorter than in the fixed continuation scenario with $\delta = 0.65$. Yet, despite the shorter expected horizon, cooperation becomes sustainable due to the stronger conditional incentives introduced by the action-dependent continuation rule.

The unconditional cooperation can be sustainable with AC strategy under the following condition:

$$\frac{1-\delta(1)}{1-\delta(2)} \cdot R \geq T. \tag{AC}$$

Hence, unconditional cooperation may be sustained as an SPE for certain parameter combinations. The intuition is as follows: a unilateral deviation to defection decreases the continuation probability (from $\delta(2)$ to $\delta(1)$), thereby reducing the expected value of future payoffs. If this reduction in the expected continuation value outweighs the benefit of receiving the temptation payoff T instead of the cooperative payoff R, then cooperation is optimal. In such cases, players may prefer to cooperate unconditionally. Figure 1 shows the combinations of $\delta(2)$ and R for which unconditional cooperation can be sustained as an SPE.

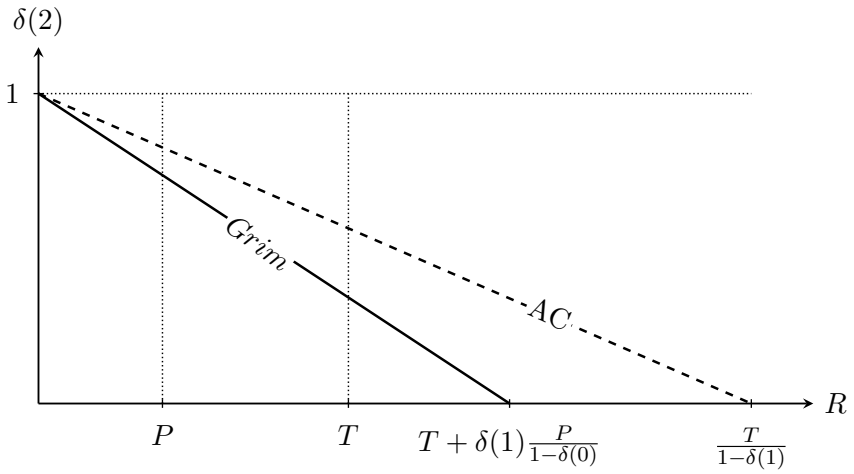


Figure 1: SPE conditions for $Grim$[11] and AC

The proofs and other conditions under which other strategies can be supported as SPE are provided in Appendix C.

4 Experiment

4.1 Experiment design

The stage game is a Prisoner's Dilemma (Table 2a) with payoffs expressed in experimental points. The payoff from mutual cooperation, R, is set equal to either 32 or 48. I consider two values for the continuation probability following mutual cooperation, $\delta(2) \in \{1/2, 2/3\}$, while holding fixed the continuation probabilities after all other action profiles at $\delta(1) = \delta(0) = 1/2$. Table 2b summarizes the continuation probabilities by action profile. Varying R and $\delta(2)$ in this way yields a 2×2 factorial design with four treatments.

Following Blonski and Spagnolo (2015), I focus on a simplified game with two strategies only: $Grim$ and AD. These strategies capture two distinct and empirically relevant forms of behavior – non-cooperative and conditionally cooperative – that may both be supported in equilibrium. Within this restricted strategy set, AD is the only noncooperative strategy that is subgame perfect. Among

[11] For Grim, condition AD also has to hold, but it does not include either $\delta(2)$ or R.

Table 2: Experimental Game

(a) Stage Game Payoffs

	C	D
C	R, R	12, 50
D	50, 12	25, 25

(b) Probability of Additional Round

	C	D
C	$\delta(2)$	1/2
D	1/2	1/2

the conditionally cooperative strategies, *Grim* is empirically relevant and the only one supportable as a subgame-perfect equilibrium. Thus, focusing on *Grim* as the cooperative strategy usually results in little loss (Dal Bó & Fréchette, 2018, 2019; Romero & Rosokha, 2018). Additionally, the focus on *Grim* and *AD* is further motivated by the empirical literature on repeated-game strategies. Strategy-elicitation and strategy-estimation exercises consistently show that behavior in the IRPD is clustered around a limited set of strategies, especially Always Defect, *Grim*, and Tit-for-Tat (Dal Bó & Fréchette, 2019; Romero & Rosokha, 2018). At the same time, among these empirically prominent pure strategies, Tit-for-Tat is not a SPE. Although the $(R = 48, \delta(2) = 2/3)$ satisfies the equilibrium conditions for unconditional cooperation I still focus on *Grim* rather than *AC*. *Grim* requires "weaker" conditions on $\delta(2)$ and R, making it the easier of the two to support. Empirical evidence suggests that cooperation is mainly conditional rather than unconditional; the wider experimental literature rarely reports unconditional cooperation. (Dal Bó & Fréchette, 2018), and estimations in Section 5.3 confirms the same pattern in this experiment: *AC* is rarely played. Thus, *Grim* and *AD* constitute the natural cooperative and noncooperative benchmark strategies. The expected-payoff matrix for the simplified game is reported in Table 3.

Table 3: Expected Payoffs

	Grim	AD
Grim	$\frac{R}{1-\delta(2)}$, $\frac{R}{1-\delta(2)}$	37, 75
AD	75, 37	50, 50

For the treatments in which cooperation can be supported, I use risk dominance and strategic uncertainty as equilibrium-selection criteria (Blonski et al., 2011; Dal Bó & Fréchette, 2011, 2018). In particular, I follow the literature and measure robustness to strategic uncertainty by the size of the basin of attraction of *AD*. Suppose player i believes that player j follows the mixed strategy $p \times Grim \otimes (1-p) \times AD$. The basin of attraction of *AD* is then the set of beliefs p for which *AD* yields a higher expected payoff than *Grim*. In the present design, this basin is the interval $[0, p^*]$,

where the cutoff p^* is given by

$$p^* = \frac{P - S}{\frac{R}{1-\delta(2)} - S - T}. \tag{1}$$

This measure also yields the ordinal risk-dominance ranking between the two strategies: if $p^* < 1/2$, then *Grim* risk dominates *AD*, and conversely.

Intuitively, as p^* grows, the set of beliefs under which *AD* is optimal expands, so cooperation should therefore decline. Existing evidence is broadly consistent with this prediction, especially in environments in which cooperation is risk dominant (Boczoń et al., 2025; Dal Bó & Fréchette, 2011, 2018). Table 4a reports the corresponding values of p^* for the treatments in which cooperation can be supported as a subgame-perfect equilibrium. In all such treatments, cooperation is risk dominant, although defection remains supportable in subgame-perfect equilibrium.

Table 4: Treatments

(a) The size basin of attraction of AD (p^*)

	$R = 32$	$R = 48$
$\delta(2) = 1/2$	not SPE	0.38
$\delta(2) = 2/3$	0.38	0.16

(b) Grim-Grim expected payoffs

	$R = 32$	$R = 48$
$\delta(2) = 1/2$	64	96
$\delta(2) = 2/3$	96	144

Figure 2 plots the treatments in relation to conditions for *Grim* and *AC* to be sustained as an equilibrium.

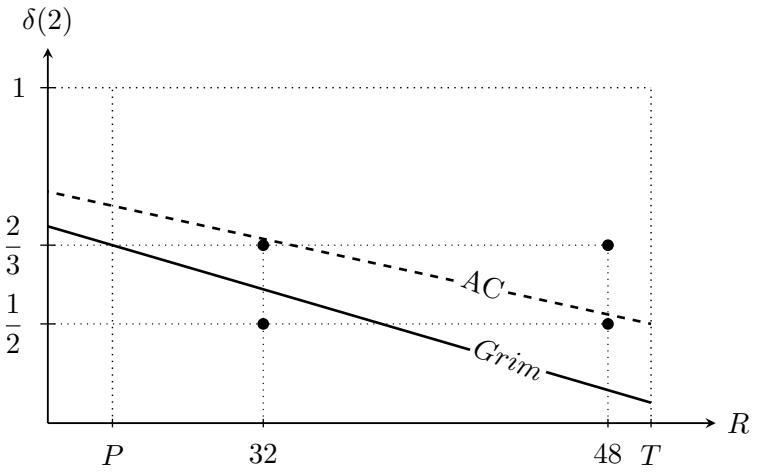


Figure 2: Experimental treatments and SPE conditions

The treatment variation is designed to change only the expected value of successful mutual cooperation. This makes it possible to isolate strategic uncertainty by varying either the payoff from mutual cooperation or the continuation probability following mutual cooperation, while holding fixed all other payoffs and continuation probabilities. As a result, the design permits a direct comparison

between the behavioral effects of increasing the reward from cooperation and those of increasing the likelihood of future interaction after joint cooperation.

Table 4b reports the expected payoff under mutual cooperation when both players choose *Grim*. Two treatments, $(R = 32, \delta(2) = 2/3)$ and $(R = 48, \delta(2) = 1/2)$, generate the same expected payoff under mutual cooperation, and therefore the identical reduced form game from Table 3. The only difference between these treatments is the mechanism through which expected cooperative payoff achieved: in one case through a higher cooperative payoff, and in the other through a higher continuation probability after mutual cooperation. This feature of the design allows for a clean comparison of how distinct incentive channels affect the use and success of cooperative strategies.

4.2 Hypotheses

The first hypothesis is motivated by subgame-perfect equilibrium selection.

Hypothesis 1. *Cooperation rate is higher in treatments where a unilateral increase in $\delta(2)$ or R renders cooperative play sustainable as a subgame-perfect equilibrium.*

This hypothesis is tested by comparing $(R = 32, \delta(2) = 1/2)$ with $(R = 48, \delta(2) = 1/2)$, and $(R = 32, \delta(2) = 1/2)$ with $(R = 32, \delta(2) = 2/3)$. In each case, the latter treatment supports cooperation in equilibrium and is therefore expected to yield higher cooperation rates through distinct channels: an increase in either R or $\delta(2)$. When cooperative equilibria are available, subjects can sustain higher levels of cooperation and attain higher payoffs. Although AD remains an equilibrium outcome, some subjects are expected to coordinate on the cooperative equilibrium.

The next set of hypotheses is based on the basin measure derived above.

Hypothesis 2. *When Grim is the risk-dominant strategy, cooperation rate increases with $\delta(2)$ or R.*

This hypothesis is tested by comparing $(R = 48, \delta(2) = 1/2)$ with $(R = 48, \delta(2) = 2/3)$, and $(R = 32, \delta(2) = 2/3)$ with $(R = 48, \delta(2) = 2/3)$. In each comparison, the latter treatment reduces the basin of attraction of AD and is therefore expected to yield higher cooperation rates. A smaller basin of attraction for AD implies a larger set of beliefs under which cooperation is optimal, which is the mechanism underlying Hypotheses 2 and 3. Additionally, an increase in either R or $\delta(2)$ makes unconditional cooperation an equilibrium strategy, which may further increase the attractiveness of cooperation. I focus on robustness to strategic uncertainty within the AD-Grim benchmark, as this provides an empirically relevant comparison for the IRPD.[12]

Hypothesis 3. *If a decrease in $\delta(2)$ is offset by an increase in R so as to keep constant the robustness to strategic uncertainty, p^*, then cooperation should remain unchanged.*

[12]The basin of attraction of AD in the AC-AD reduced game is $[0, 0.37]$, implying $p^* = 0.37$. Thus, the treatment with $R = 48$ and $\delta(2) = 2/3$ is as robust to strategic uncertainty under the AD-AC framework, with *AC* being more "risky" strategy comparing to *Grim*.

This hypothesis is tested by comparing $(R = 48, \delta(2) = 1/2)$ with $(R = 32, \delta(2) = 2/3)$. These two treatments keep the basin of attraction of AD, p^*, constant by generating the same game matrix in the simplified strategy space. Under equilibrium selection based on robustness to strategic uncertainty, cooperation should therefore be similar across the two treatments.

4.3 Session details

Sessions were conducted in a university experimental laboratory. At the start of each session, participants read written instructions individually, after which the instructions were summarized publicly to ensure common understanding. Participants completed a comprehension quiz to verify their understanding of the payoff structure and were not permitted to proceed until all items were answered correctly. All decisions were made privately. Participants interacted anonymously through computer terminals and played a sequence of indefinitely repeated prisoner's dilemma (IRPD) games. At the end of each supergame, subjects were randomly rematched with a new partner. Sessions continued until the first supergame to end after 30 minutes of play, which determined the end of the session. Treatment parameters were fixed within session, so that each session corresponded to a single treatment condition. Five sessions were conducted for each treatment. The experiment was implemented in oTree (Chen et al., 2016). Within each supergame, participants observed the full history of play. Sample instructions from one session are reported in Appendix E.

To accommodate heterogeneity in realized supergame length induced by the action-dependent continuation rule while maintaining comparability across sessions, I implemented the Block Random Termination (BRT) procedure of Fréchette and Yuksel (2017). Participants played each supergame in blocks of three rounds. During a block, participants were not informed whether the match would continue beyond that block or at which round it had terminated. At the end of each three-round block, an additional block was played whenever the supergame had not terminated for *at least one pair* in the session. This procedure synchronized continuation more closely across sessions and facilitated tractable data analysis while preserving the random termination structure.[13]

Although prior work on indefinitely repeated prisoner's dilemma games finds limited evidence that risk aversion is an important determinant of behavior (Dal Bó & Fréchette, 2018), risk preferences may nonetheless affect the interpretation of Hypothesis 3.[14] To account for this possibility, I elicited subjects' risk preferences at the end of each session using the Bomb Risk Elicitation Task (Crosetto & Filippin, 2013; Holzmeister & Pfurtscheller, 2016). The instructions for this task were provided only after subjects had completed their final supergame, in order to avoid priming effects. The results indicate that risk attitudes do not affect cooperation across treatments; further details are

[13] Supergame termination was determined by a pseudo-random number generator. To limit variation arising from a randomization, each treatment condition was exposed to the same set of random seeds. Termination was determined by a single die roll per round, common to all pairs in a session.

[14] Under a non-risk-neutral utility function over monetary payoffs, $U(\cdot)$, the expected payoffs under mutual cooperation in the two treatments considered in Hypothesis 3 are not equal: $3U(32) \neq 2U(48)$.

provided in Appendix B. Following the risk elicitation task, subjects completed a short demographic questionnaire and responded to an open-ended question regarding the considerations that guided their behavior during the experiment.

Table 5 summarizes the session-level information. In total, the experiment involved 224 participants, each of whom made decisions in 34.45 supergames on average. Sessions lasted for about 70 minutes. Participants were paid according to their cumulative earnings from all rounds of one randomly selected supergame, which preserves incentive compatibility (Azrieli et al., 2018).[15] In addition to a $5 show-up fee, participants earned $18.97 on average from their experimental decisions. Total earnings ranged from $6.80 to $53.15.

Table 5: Session Summary Information

	Subjects	Rounds			Matches			Earnings ($)		
Treatment	per session	Play[a]	Count[b]	Diff	Avg	Min	Max	Avg	Min	Max
$\delta(2) = 1/2,\ R = 32$	8, 12, 12, 12, 12	3.32	1.86	1.46	37.2	32	42	18.82	8.75	45.20
$\delta(2) = 1/2,\ R = 48$	12, 12, 12, 10, 14	3.34	1.88	1.47	35.6	22	43	20.63	6.80	44.90
$\delta(2) = 2/3,\ R = 32$	10, 12, 12, 12, 10	3.99	2.05	1.93	30.4	19	37	17.55	6.80	48.20
$\delta(2) = 2/3,\ R = 48$	8, 12, 12, 12, 8	4.09	2.17	1.92	34.6	31	38	18.73	8.00	53.15

[a] Average rounds played, including rounds that may not count towards earnings in BRT.
[b] Average rounds counted towards final payoff.

Each supergame consisted of an average of 3.67 rounds, of which only 1.99 rounds, on average, contributed to participants' potential final payoff. This discrepancy reflects the Block Random Termination procedure implemented in the experiment. Table 5 reports both the total number of rounds played and the number of payoff-relevant rounds across treatment conditions. Notably, treatments with a contingent continuation rule ($\delta(2) = 2/3$) exhibit a larger share of rounds that did not count toward final earnings. One possible concern is that this feature may have indirectly encouraged cooperation. Because cooperative pairs were more likely to trigger additional blocks, non-cooperative subjects could end up playing more rounds that ultimately did not affect payoffs. To the extent that these rounds still required time and effort, subjects may have been nudged toward cooperation in order to reduce exposure to such unproductive interactions, potentially confounding the effect of contingent continuation.

Two considerations address this concern. First, the differences in proportions are modest: 56% of rounds were payoff-relevant under fixed continuation probability, compared with 52% under contingent continuation. These difference appear too small, on average, to account for substantial behavioral variation. Second, because only one supergame was randomly selected for payment,

[15] Sherstyuk et al. (2013) compare payment based on all rounds with payment based only on the last round in indefinitely repeated games and find no significant behavioral differences conditional on the payment rule. However, as noted by Fréchette et al. (2024), there is no direct experimental evidence on whether behavior differs when payment is based on one rather than multiple supergames.

participants had no incentive to strategically prolong play in order to reduce the number of payoff-irrelevant rounds in the session. Moreover, the analysis below shows that exposure to payoff-irrelevant blocks does not explain the observed cooperation patterns.

5 Results

5.1 Overall description

Table 6 reports cooperation rates by treatment, distinguishing between cooperation in the first round of each supergame (left panel) and average cooperation across all rounds (right panel). I report first-round cooperation separately because it is the only decision every subject makes in every supergame before any history with the current partner exists. Averages over all rounds also mix in later rounds, whose number differs across treatments and in which play responds to the partner's past actions. Under contingent continuation this matters more than usual: cooperative pairs are more likely to survive to later rounds, so all-round averages give extra weight to cooperative matches.

Three main patterns emerge from the treatment comparisons. First, Cooperation rates are higher when a unilateral increase in either R or $\delta(2)$ makes cooperative play sustainable as a SPE, which supports the Hypothesis 1. I first compare $(R = 32, \delta(2) = 1/2)$ with $(R = 48, \delta(2) = 1/2)$. Cooperation is significantly higher in the latter treatment, indicating that increasing the reward from mutual cooperation raises cooperation when the increase makes cooperative play sustainable in equilibrium. I then compare $(R = 32, \delta(2) = 1/2)$ with $(R = 32, \delta(2) = 2/3)$. Cooperation is again significantly higher in the latter treatment, indicating that increasing the continuation probability following mutual cooperation raises cooperation when it makes cooperative play sustainable. Both comparisons are statistically significant for first-round cooperation and average cooperation across all rounds.

Second, once conditionally cooperative play is sustainable, further increase in R or $\delta(2)$ does not significantly increase cooperation rates, so Hypothesis 2 is not supported. I first compare $(R = 48, \delta(2) = 1/2)$ with $(R = 48, \delta(2) = 2/3)$. Although the latter treatment has a smaller basin of attraction for AD and higher expected payoff of cooperative play, increasing the continuation probability does not produce a statistically significant increase in cooperation. We then compare $(R = 32, \delta(2) = 2/3)$ with $(R = 48, \delta(2) = 2/3)$. The latter treatment yields a higher point estimate of cooperation, but the difference is not statistically significant.

Third, cooperation rate does not differ significantly across treatments with the same robustness to strategic uncertainty. I compare $(R = 48, \delta(2) = 1/2)$ with $(R = 32, \delta(2) = 2/3)$. These treatments generate the same value of p^* but differ in whether incentives are strengthened through the cooperative reward or the continuation probability. The treatment with the higher cooperative reward yields a larger point estimate of cooperation, but the difference is not statistically significant for either first-round cooperation ($p = 0.3207$) or average cooperation across all rounds ($p = 0.5411$).

We therefore cannot reject the null hypothesis that cooperation is the same across the two treatments. This finding is consistent with Hypothesis 3, although the absence of a significant difference does not establish statistical equivalence.

The qualitative conclusions are unchanged under wild-cluster-bootstrap inference, player-level clustering, and rank-sum tests. Additional results are reported in Appendix A.

Table 6: Cooperation Rates in All Repeated Games

	First round			All rounds		
$\delta(2)\backslash R$	32		48	32		48
1/2	0.06	$<^{***}$	0.47	0.06	$<^{***}$	0.35
	$\wedge^{***}$	$\angle$	?|	$\wedge^{***}$	$\angle$	?|
2/3	0.36	$<$	0.44	0.30	$<$	0.35

Notes: Statistical significance is estimated with probit regressions with a dummy variable for one of the relevant treatments. Standard errors are clustered by session.

To better understand how these treatment effects arise, it is useful to focus on behavior at the very first supergame (Table 7). Two main patterns emerge. First, in the first round of the first supergame, subjects respond strongly to increases in the reward from mutual cooperation, whereas an increase in the continuation probability has no statistically significant effect. This suggests that, in the initial interaction, subjects primarily condition their behavior on contemporaneous payoffs and do not fully internalize the dynamic incentives generated by a higher continuation probability following joint cooperation.

Second, when behavior is aggregated over all rounds of the first supergame, cooperation is broadly consistent with the comparative statics implied by the hypotheses. Cooperation increases with the reward from mutual cooperation for both levels of the continuation probability. Moreover, in the low-reward treatments ($R = 32$), a higher continuation probability leads to more cooperation. By contrast, when the reward is high ($R = 48$), changing the continuation probability does not have a statistically significant effect on cooperation. Thus, while a higher continuation probability promotes cooperation when it makes cooperation an equilibrium outcome, its effect appears muted once the reward from mutual cooperation is sufficiently high and cooperation can be supported as an equilibrium outcome under both continuation rules. Almost all[16] results are unchanged under additional procedures reported in Appendix A.

[16]The only not robust finding is the effect of continuation probability when $R = 32$ in all rounds of the first repeated game.

Table 7: Cooperation Rates in First Repeated Game

	First round			All rounds		
$\delta(2)\backslash R$	32		48	32		48
1/2	0.38	$<^{***}$	0.62	0.16	$<^{***}$	0.53
		$\lesssim^{***}$		$\wedge^{**}$	$\lesssim^{***}$	
2/3	0.38	$<^{***}$	0.65	0.31	$<^{***}$	0.49

Notes: Statistical significance is estimated with probit regressions with a dummy variable for one of the relevant treatments. Standard errors are clustered by session.

5.2 Evolution of behavior

Previous work shows that learning and the evolution of behavior within the experiment can have important effects on cooperation rates (Dal Bó & Fréchette, 2011). Consistent with this literature, the patterns in my data differ depending on whether one considers all supergames or only the first supergame. Figure 3 plots average first-round cooperation by match for each treatment, while Figure 4 reports the corresponding first-round cooperation rates by session and treatment. As shown above, subjects begin with higher cooperation under the higher reward, whereas differences across continuation probabilities are comparatively small.

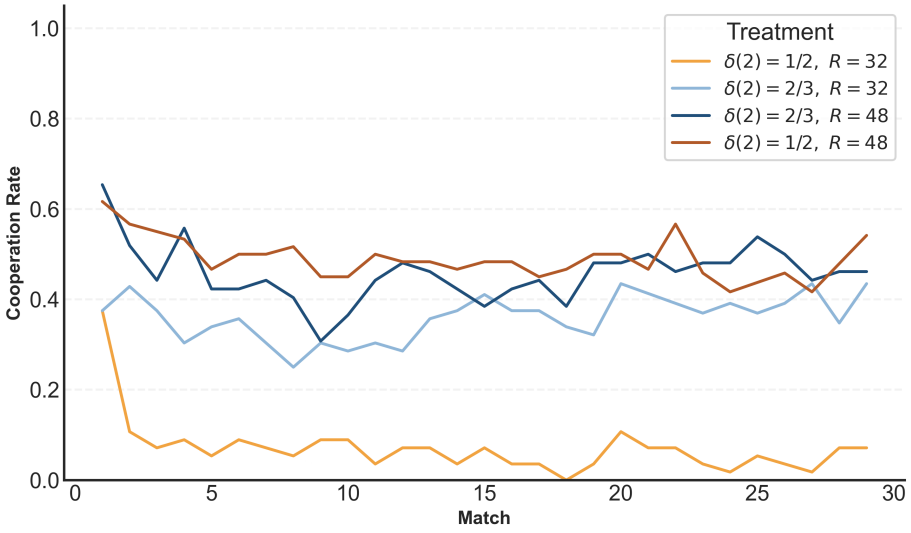
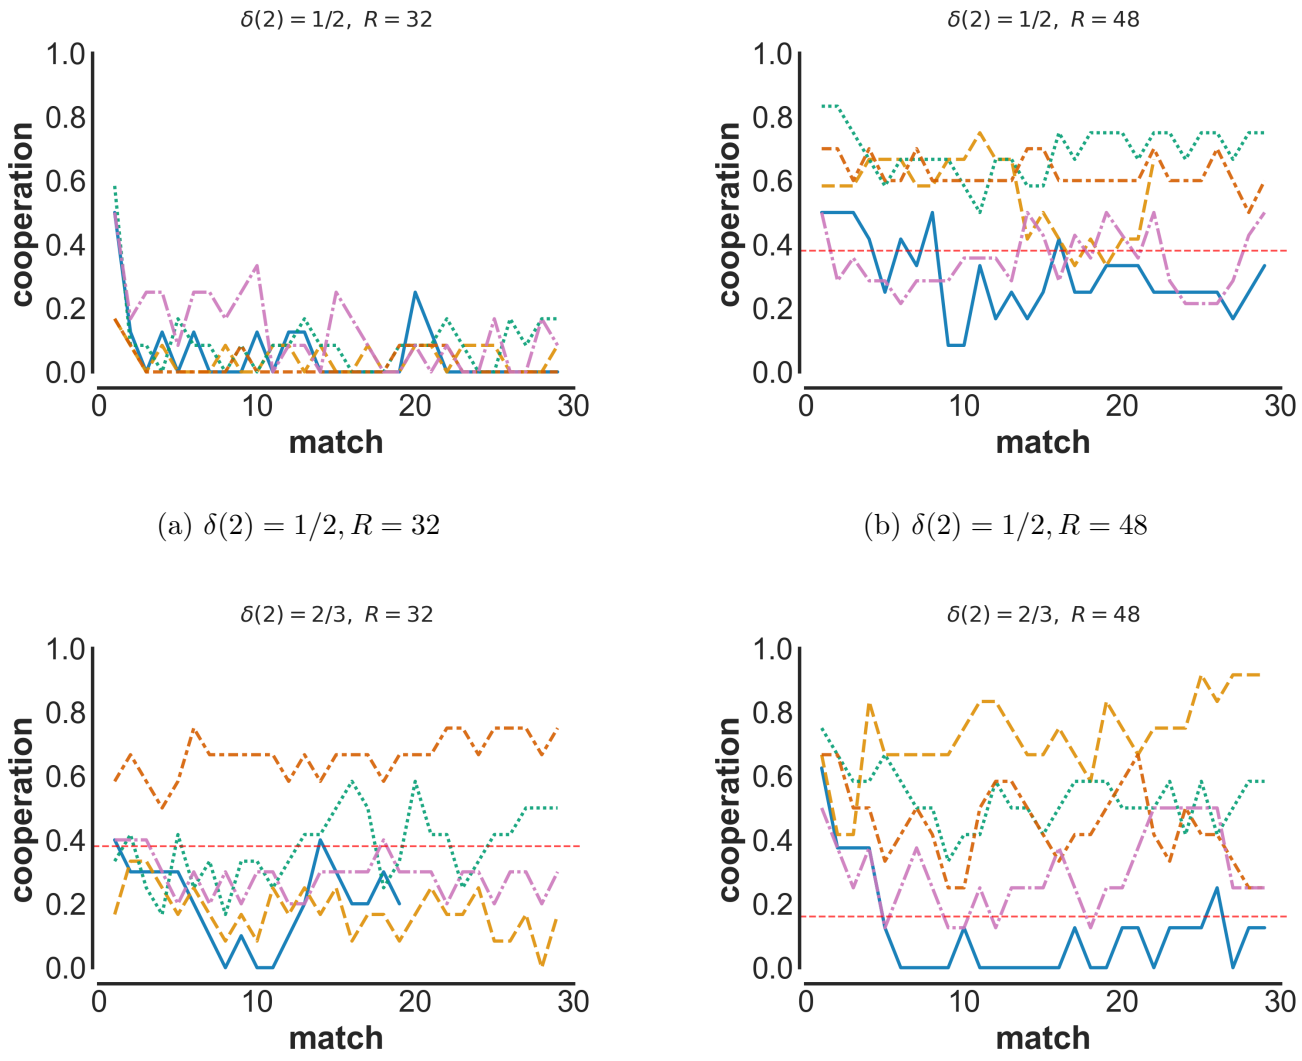


Figure 3: Evolution of cooperation for first rounds

Notes: Each line plots the average first-round cooperation rate across all sessions within a treatment, by supergame number. Treatment parameters are indicated in the legend.

Subjects in treatment in which cooperation is not a subgame-perfect equilibrium action learn relatively quickly not to cooperate. In the remaining treatments, however, cooperation declines



(a) $\delta(2) = 1/2, R = 32$

(b) $\delta(2) = 1/2, R = 48$

(c) $\delta(2) = 2/3, R = 32$

(d) $\delta(2) = 2/3, R = 48$

Figure 4: Evolution of Cooperation by Treatment and Session (First Rounds)

Notes: Each line denotes the first-round cooperation rate in a different session within a treatment. Horizontal dotted lines denote the basin-of-attraction threshold p^* separating Always Defect and *Grim* strategies. No threshold is displayed in panel (a) because cooperation cannot be supported as a subgame-perfect equilibrium in that treatment.

during roughly the first ten supergames and then becomes comparatively stable. This makes the within-session dynamics particularly important. As Figure 4 shows, there is substantial heterogeneity across sessions within the same treatment.

Two patterns are particularly noteworthy. First, in the high-reward treatment with contingent continuation probability (Figure 4d), sessions begin from broadly similar levels but diverge sharply over the course of the experiment. One possible explanation is that sessions differ in the continuation histories subjects experience. Second, in the treatment with $R = 48$ and $\delta(2) = 1/2$ (Figure 4b), session-level cooperation tends to remain either mostly above or mostly below the basin-of-attraction threshold for Always Defect, p^*. Since cooperation is optimal when subjects expect others to coop-

erate with probability above p^*, and defection is optimal otherwise, this pattern is consistent with subjects conditioning their behavior on beliefs about others' likely cooperation. If subjects form reasonably accurate beliefs, as suggested by Aoyagi et al. (2024), then lower expected cooperation should lead to lower realized cooperation. A similar, though weaker, pattern appears in the treatment with $R = 32$ and $\delta(2) = 2/3$ (Figure 4c). Because higher rewards generate higher initial cooperation, they also increase the likelihood that a session remains above p^*, thereby sustaining higher cooperation over time.

To examine the role of prior experience more directly, I study how cooperation in a given supergame depends on the previous partner's behavior and on the length of the preceding supergame. Table 8 reports average marginal effects from probit regressions estimated on all matches except the first. First, cooperation in later supergames is strongly related to first-round cooperation in the initial supergame. Since higher rewards also produce higher first-round cooperation at the outset, this points to a stronger role for the reward parameter than for the continuation rule in shaping subsequent cooperation. Second, under contingent continuation probability, more payoff relevant rounds in previous matches are associated with higher cooperation in the next supergame. This suggests that subjects learn from experiencing the consequences of potentially longer interactions and incorporate that experience into subsequent first-round decisions. It also helps account for the substantial session heterogeneity under contingent continuation: when random realizations of cooperation affect the experienced duration of play, early differences can propagate over time.

More generally, across all treatments, subjects are more likely to begin the next supergame by cooperating when their previous partner cooperated in the first round of the preceding supergame. This marginal effect is largest in the treatment that exhibits the greatest inter-session heterogeneity, $\delta(2) = 2/3; R = 48$. This pattern is consistent with belief updating (Dal Bó & Fréchette, 2011). A partner who opened with cooperation is evidence that cooperative strategies are common in the session. That evidence moves a subject's belief further above p^* and raises the expected return to opening the next match with cooperation.

Finally, payoff-irrelevant blocks may convey information about others' behavior, for example that cooperation elsewhere in the session prolonged interaction. If subjects used this information to update their beliefs, having played a full irrelevant block in the previous match should raise cooperation in the next one, confounding the effect of contingent continuation. Table 8 shows no such effect: the coefficient is negative at $R = 32$ and close to zero at $R = 48$, and neither is significant.

5.3 Strategies

This section examines the strategies used by participants. To do so, I employ the Strategy Frequency Estimation Method (SFEM) introduced by Dal Bó and Fréchette (2011), which uses maximum likelihood to infer the prevalence of different strategies from the observed round-by-round cooperation decisions. The estimation procedure selects the distribution of strategies and the associated error

Table 8: Marginal Effects of Factors Affecting the Evolution of Behavior

	$\delta(2) = 1/2$		$\delta(2) = 2/3$	
	$R = 32$	$R = 48$	$R = 32$	$R = 48$
Subject cooperated in round 1 of match 1	0.085***	0.427***	0.301***	0.255***
	(0.024)	(0.096)	(0.082)	(0.079)
Partner cooperated in round 1 of prev. match	0.032***	0.125***	0.160***	0.236***
	(0.008)	(0.021)	(0.025)	(0.078)
Relevant rounds in previous match	0.004	-0.004	0.025***	0.014*
	(0.004)	(0.009)	(0.008)	(0.008)
Played full irrelevant block in previous match			-0.088	0.006
			(0.057)	(0.059)
Match	-0.001	-0.002	0.002	-0.000
	(0.001)	(0.002)	(0.003)	(0.001)
Observations	2048.000	2078.000	1656.000	1728.000

Notes: Clustered SEs at session level
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

probability that jointly maximize the likelihood of the observed sequence of play. For the present analysis, I focus on a core set of eleven strategies, following Dal Bó and Fréchette (2019) and Fudenberg et al. (2012). Table 9 lists the strategies included in the estimation and provides their descriptions. The estimation is implemented using the code developed by Romero and Rosokha (2018) and Rosokha and Wei (2024).

Table 9: Descriptions of the 11 Strategies Considered

Strategy	**Abbreviation**	**Description**
Always Cooperate	AC	Always play C.
Tit-for-Tat	TFT	Play C unless partner played D last round.
Tit-for-2-Tats	TF2T	Play C unless partner played D in both of the last 2 rounds.
Tit-for-3-Tats	TF3T	Play C unless partner played D in all of the last 3 rounds.
2-Tits-for-1-Tat	2TFT	Play C unless partner played D in either of the last 2 rounds (2 rounds of punishment if partner plays D).
2-Tits-for-2-Tats	2TF2T	Play C unless partner played 2 consecutive Ds in the last 3 rounds (2 rounds of punishment if partner plays D twice in a row).
Grim	Grim	Play C until either player plays D, then play D forever.
Lenient Grim 2	Grim2	Play C until 2 consecutive rounds occur in which either player played D, then play D forever.
Lenient Grim 3	Grim3	Play C until 3 consecutive rounds occur in which either player played D, then play D forever.
Suspicious Tit-for-Tat	STFT	Play D unless partner played C last round.
Always Defect	AD	Always play D.

I report the estimation results for supergames after the fifth in Table 10, since the initial supergames likely reflect subjects' experimentation and learning about available strategies. Table 11

reports the estimated pairwise differences in strategy shares across treatments.

Several patterns emerge from the estimates. First, the frequency of AD does not differ across treatments once cooperation is supportable in both compared cells. Beyond differences in the overall prevalence of AD, the treatment variation appears to operate mainly through changes in the composition of conditionally cooperative strategies. This suggests that subjects respond to the strategic environment primarily by adjusting the severity of contingent punishment.

Second, *Grim* is not used significantly by subjects (point estimates reach 0.11 under $\delta(2) = 2/3$ but are not statistically significant). Instead, there is significant use of other conditionally cooperative strategies that begin with cooperation, in particular TFT and 2TFT, which broadly consistent the previous literature with $\delta = 1/2$. TFT use varies significantly across treatments, with the lowest frequencies arising at the two extremes: when cooperation is not a subgame-perfect equilibrium and when cooperation is most robust to strategic uncertainty. This pattern is consistent with the fact that conditional cooperation is never a best response in the former case, whereas in the latter case subjects rely more heavily on conditionally cooperative strategies with harsher punishment, such as 2TFT.[17]

Third, in three of the four treatments the estimated frequency of STFT exceeds that of TFT. This contrasts with prior findings in the strategy estimation and strategy elicitation literature. One possible explanation is that session-level variation in the evolution of play generated more unstable beliefs, making subjects less confident in the reliability of their matched partners. This, in turn, may have favored the use of more cautious strategies such as STFT. At the same time, there is no significant difference in the use of STFT across treatments.

Fourth, there is a marginally statistically significant share of unconditional cooperation in treatments with $\delta(2) = 48$, but these frequencies are not economically meaningful. Additionally, AC can be sustained as an equilibrium in one of these treatments but not in the other, yet there is no statistically significant difference in the frequency with which subjects employ this strategy. This provides evidence that, even when unconditional cooperation can be sustained as an equilibrium, subjects rarely employ an unconditionally cooperative strategy.

Finally, there are no differences in strategy use between $(R = 48, \delta(2) = 1/2)$ and $(R = 32, \delta(2) = 2/3)$, even though these treatments differ in their underlying parameters. Since they share the same simplified strategy-game matrix and the same degree of robustness to strategic uncertainty, this finding is consistent with (but not establishing) substitution at the aggregate level, even though considerable heterogeneity across sessions persists within each treatment.

The near-zero estimated shares of TF2T, TF3T, 2TF2T, Grim2, and Grim3 further indicate that behavior is concentrated in a small set of relatively simple strategies, implying that the empirically relevant margins are between unconditional defection, cautious reciprocity, standard reciprocity, and

[17] In matches as short as those observed here, several punishment strategies prescribe the same actions. After a defection in round 1, *Grim* and 2TFT both defect through round 3 and separate only if play continues and the opponent returns to cooperation due to restart effect with the new block (Fréchette & Yuksel, 2017). The shares of these strategies are therefore less precisely identified than the shares of AD and of conditional cooperation as a whole.

harsher short-memory punishment.

Table 10: Estimation of Strategies Used

	$\delta(2)=1/2$		$\delta(2)=2/3$	
	$R=32$	$R=48$	$R=32$	$R=48$
β	0.96	0.90	0.91	0.89
AD	0.76***	0.48***	0.37***	0.38***
	(0.06)	(0.07)	(0.06)	(0.07)
STFT	0.22***	0.19***	0.12***	0.20***
	(0.06)	(0.05)	(0.04)	(0.05)
2TFT	0.00	0.13***	0.13*	0.23**
	(0.0)	(0.05)	(0.1)	(0.1)
TFT	0.00	0.18***	0.24***	0.04
	(0.0)	(0.06)	(0.06)	(0.03)
Grim	0.00	0.00	0.11	0.11
	(0.0)	(0.0)	(0.1)	(0.09)
AC	0.02	0.00	0.03*	0.04*
	(0.02)	(0.0)	(0.02)	(0.03)

Notes: Strategies with frequency ≤ 0.01 are not included. Participant-block nonparametric bootstrap standard errors are in parentheses. β is the probability that a subject implements the action prescribed by the strategy. Supergames after fifth are used for estimation. One-sided p-values: *** $p<0.01$, ** $p<0.05$, * $p<0.10$.

Table 11: Pairwise Differences in Strategy Shares

	ΔR (32 vs 48)		$\Delta\delta$ (1/2 vs 2/3)		Diagonal
	$\delta=1/2$	$\delta=2/3$	$R=32$	$R=48$	($\delta=2/3, R=32$) vs ($\delta=1/2, R=48$)
AD	-0.28***	0.01	-0.39***	-0.10	0.11
	(0.09)	(0.09)	(0.08)	(0.10)	(0.09)
STFT	-0.03	0.08	-0.10	0.01	0.07
	(0.08)	(0.06)	(0.07)	(0.07)	(0.06)
2TFT	0.13***	0.10	0.13	0.10	0.00
	(0.05)	(0.14)	(0.10)	(0.11)	(0.11)
TFT	0.18***	-0.20***	0.24***	-0.14**	-0.06
	(0.06)	(0.07)	(0.06)	(0.07)	(0.08)
Grim	0.00	0.00	0.11	0.11	-0.11
	(0.00)	(0.13)	(0.10)	(0.09)	(0.10)
AC	-0.02	0.01	0.01	0.04	-0.03
	(0.02)	(0.04)	(0.03)	(0.03)	(0.02)

Notes: The table presents the pairwise differences in used strategies between treatments. For example, column 3 ($\Delta\delta, R=32$) shows the differences in the share of used strategies between treatment ($\delta(2)=1/2, R=32$) and ($\delta(2)=2/3, R=32$). Strategies with frequency ≤ 0.01 are not included. Standard errors in parentheses.
*** $p<0.01$, ** $p<0.05$, * $p<0.10$.

6 Conclusion

This paper introduces an indefinitely repeated prisoner's dilemma in which cooperation raises the probability that the interaction continues. I ask whether incentives operating through the survival of the relationship can substitute for incentives operating through payoffs. Three findings emerge. Contingent continuation raises cooperation rates where it makes cooperation supportable in equilibrium, but produces no detectable increase where cooperation is already supportable. Payoff incentives raise cooperation immediately, whereas theoretically equivalent continuation incentives take hold only with experience.

Why do payoff incentives bite earlier? The natural explanation is representational rather than experiential. The reward from mutual cooperation is displayed directly in the payoff table. The continuation benefit is presented as a rule about die rolls, and appreciating its value requires converting that rule into an expected horizon and the horizon into an expected payoff stream. The timing in the data supports this account: the asymmetry is present in the very first round of the very first supergame, before any subject has experienced a longer match or observed a higher realized payoff, so it cannot be driven by differential learning from outcomes. So does the subsequent convergence, since experience with the continuation rule supplies through observation what was initially available only through computation.

An alternative explanation is that subjects misperceive probabilities. Under prospect theory (Kahneman & Tversky, 1979), moderate to high probabilities are underweighted, so the subjective distance between deltas may be smaller than the objective one, while the reward enters utility directly and is not subject to the same compression. The two accounts differ in a testable way. Probability weighting is a stable feature of preferences and predicts a persistent asymmetry; a computational account predicts that the asymmetry attenuates as subjects gain experience with the continuation rule. The observed convergence therefore favors the latter, although the two are not mutually exclusive and the design was not built to separate them.

One theoretical implication of the environment is deliberately left untested. When $\delta(1)$ is greater than $\delta(0)$, unilateral cooperation against a defector lengthens the expected relationship, and for some parameter values this makes permanent mutual defection unsustainable in equilibrium. The experiment holds those $\delta(0)$ and $\delta(1)$ fixed so Always Defect remains an equilibrium in all four treatments. Fixing the off-path continuation probabilities is what allows the treatments to isolate the value of joint cooperation but it means the paper tests the substitutability prediction and not the unraveling one.

The results carry a message for the settings that motivate the paper. The canonical IRPD is nested here as the special case in which the continuation probabilities coincide, and its central regularities extend to environments with an endogenous horizon. And for institutions that tie the survival of a relationship to conduct, from cartel governance to alliance design, the equivalence of survival based and payoff based discipline is a property of experienced participants. In setting where

parties are new to the environment, immediate rewards take hold faster, and the two instruments should not be treated as interchangeable from the outset.

References

Abreu, D., Pearce, D., & Stacchetti, E. (1990). Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica: Journal of the Econometric Society*, 1041–1063.

Anderson, L. R., Hicks, R. L., & Turscak, A. (2024). Cooperation in public goods experiments with random and finite stopping rules. *Journal of the Economic Science Association*, 1–26.

Aoyagi, M., Bhaskar, V., & Fréchette, G. R. (2019). The impact of monitoring in infinitely repeated games: Perfect, public, and private. *American Economic Journal: Microeconomics*, *11*(1), 1–43.

Aoyagi, M., Fréchette, G. R., & Yuksel, S. (2024). Beliefs in repeated games: An experiment. *American Economic Review*, *114*(12), 3944–3975.

Axelrod, R. (1984). *The evolution of cooperation*. Basic Books.

Azrieli, Y., Chambers, C. P., & Healy, P. J. (2018). Incentives in experiments: A theoretical analysis. *Journal of Political Economy*, *126*(4), 1472–1503.

Battaglini, M., Nunnari, S., & Palfrey, T. R. (2016). The dynamic free rider problem: A laboratory study. *American Economic Journal: Microeconomics*, *8*(4), 268–308.

Bigoni, M., Casari, M., Skrzypacz, A., & Spagnolo, G. (2015). Time horizon and cooperation in continuous time. *Econometrica*, *83*(2), 587–616.

Blonski, M., Ockenfels, P., & Spagnolo, G. (2011). Equilibrium selection in the repeated prisoner's dilemma: Axiomatic approach and experimental evidence. *American Economic Journal: Microeconomics*, *3*(3), 164–192.

Blonski, M., & Spagnolo, G. (2015). Prisoners' other dilemma. *International Journal of Game Theory*, *44*, 61–81.

Boczoń, M., Vespa, E., Weidman, T., & Wilson, A. J. (2025). Testing models of strategic uncertainty: Equilibrium selection in repeated games. *Journal of the European Economic Association*, *23*(2), 784–814.

Chen, D. L., Schonger, M., & Wickens, C. (2016). Otree—an open-source platform for laboratory, online, and field experiments. *Journal of Behavioral and Experimental Finance*, *9*, 88–97.

Crosetto, P., & Filippin, A. (2013). The "bomb" risk elicitation task. *Journal of risk and uncertainty*, *47*, 31–65.

Dal Bó, P. (2005). Cooperation under the shadow of the future: Experimental evidence from infinitely repeated games. *American economic review*, *95*(5), 1591–1604.

Dal Bó, P., & Fréchette, G. R. (2011). The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, *101*(1), 411–429.

Dal Bó, P., & Fréchette, G. R. (2018). On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature*, *56*(1), 60–114.

Dal Bó, P., & Fréchette, G. R. (2019). Strategy choice in the infinitely repeated prisoner's dilemma. *American Economic Review*, *109*(11), 3929–52.

Fréchette, G. R., Vespa, E., & Yüksel, S. (2024, December). Repeated games [Prepared for the *Handbook of Experimental Economics Methods*. Forthcoming]. In *Handbook of experimental economics methods*.

Fréchette, G. R., & Yuksel, S. (2017). Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination. *Experimental Economics*, *20*, 279–308.

Friedman, J. W. (1971). A non-cooperative equilibrium for supergames. *The Review of Economic Studies*, *38*(1), 1–12.

Fudenberg, D., & Maskin, E. (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, *54*(3), 533–554.

Fudenberg, D., Rand, D. G., & Dreber, A. (2012). Slow to anger and fast to forgive: Cooperation in an uncertain world. *American Economic Review*, *102*(2), 720–749.

Fujiwara-Greve, T., & Okuno-Fujiwara, M. (2009). Voluntarily separable repeated Prisoner's Dilemma. *Review of Economic Studies*, *76*(3), 993–1021.

Fujiwara-Greve, T., Okuno-Fujiwara, M., & Suzuki, N. (2012). Voluntarily separable repeated Prisoner's Dilemma with reference letters. *Games and Economic Behavior*, *74*(2), 504–516.

Ghosh, P., & Ray, D. (1996). Cooperation in community interaction without information flows. *Review of Economic Studies*, *63*(3), 491–519.

Gill, D., & Rosokha, Y. (2024). Beliefs, learning, and personality in the indefinitely repeated prisoner's dilemma. *American Economic Journal: Microeconomics*, *16*(3), 259–283.

Holzmeister, F., & Pfurtscheller, A. (2016). Otree: The "bomb" risk elicitation task. *Journal of Behavioral and Experimental Finance*, *10*, 105–108.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*(2), 263–291.

Kamei, K. (2019). Cooperation and endogenous repetition in an infinitely repeated social dilemma. *International Journal of Game Theory*, *48*(3), 797–834.

Kline, P., & Santos, A. (2012). A score based approach to wild bootstrap inference. *Journal of Econometric Methods*, *1*(1), 23–41.

Kloosterman, A. (2020). Cooperation in stochastic games: A prisoner's dilemma experiment. *Experimental Economics*, *23*(2), 447–467.

Kochov, A., & Song, Y. (2023). Intertemporal hedging and trade in repeated games with recursive utility. *Econometrica*, *91*(6), 2333–2369.

Lee, N. (2020, October). *An experiment: Voluntary separation in indefinitely repeated prisoner's dilemma game* (tech. rep.). SSRN Working Paper.

Lugovskyy, V., Puzzello, D., Sorensen, A., Walker, J., & Williams, A. (2017). An experimental study of finitely and infinitely repeated linear public goods games. *Games and Economic Behavior*, *102*, 286–302.

Lugovskyy, V., Puzzello, D., & Walker, J. (2020). *On cooperation in finitely and indefinitely repeated prisoner's dilemma games* [Working paper, Indiana University].

Mailath, G. J., & Samuelson, L. (2006). *Repeated games and reputations: Long-run relationships*. Oxford university press.

Mengel, F., Orlandi, L., & Weidenholzer, S. (2022). Match length realization and cooperation in indefinitely repeated games. *Journal of Economic Theory*, *200*, 105416.

Proto, E., Rustichini, A., & Sofianos, A. (2019). Intelligence, personality, and gains from cooperation in repeated interactions. *Journal of Political Economy*, *127*(3), 1351–1390.

Rob, R., & Yang, H. (2010). Long-term relationships as safeguards. *Economic Theory*, *43*(2), 143–166.

Romero, J., & Rosokha, Y. (2018). Constructing strategies in the indefinitely repeated prisoner's dilemma game. *European Economic Review*, *104*, 185–219.

Rosokha, Y., & Wei, C. (2024). Cooperation in queueing systems. *Management Science*, *70*(11), 7597–7616.

Roth, A. E., & Murnighan, J. K. (1978). Equilibrium behavior and repeated play of the prisoner's dilemma. *Journal of Mathematical psychology*, *17*(2), 189–198.

Sekiguchi, T., & Wakai, K. (2016). *Repeated games with recursive utility: Cournot duopoly under gain/loss asymmetry* (tech. rep.).

Sherstyuk, K., Tarui, N., & Saijo, T. (2013). Payment schemes in infinite-horizon experimental games. *Experimental Economics*, *16*, 125–153.

Thomas, E. A. C., & Feldman, M. W. (1988). Behavior-dependent contexts for repeated plays of the prisoner's dilemma. *Journal of Conflict Resolution*, *32*(4), 699–726.

Vespa, E. (2020). An experimental investigation of cooperation in the dynamic common pool game. *International Economic Review*, *61*(1), 417–440.

Vespa, E., & Wilson, A. J. (2019). Experimenting with the transition rule in dynamic games. *Quantitative Economics*, *10*(4), 1825–1849.

Wilson, A. J., & Wu, H. (2017). At-will relationships: How an option to walk away affects cooperation and efficiency. *Games and Economic Behavior*, *102*, 487–507.

A Additional Analysis of Cooperation Rates

The cooperation results in Section 5 use probit regressions with standard errors clustered at the session level, an approximation that is asymptotic in the number of clusters, of which there are only five per treatment. This appendix re-examines every comparison in Tables 6 and 7 in two ways. The first keeps the decision-level probit contrast and changes only the inference: (i) a score wild cluster bootstrap with session clusters (Kline & Santos, 2012), and (ii) standard errors clustered on the subject. The second replaces the contrast with nonparametric comparisons of aggregated outcomes: (iii) a two-sided Wilcoxon rank-sum test on session means, which weights sessions equally regardless of size and cannot produce a p-value below $2/\binom{10}{5} = 0.008$ with five sessions per arm.

Table A1: Cooperation Rates in All Repeated Games — alternative inference

	First round			All rounds		
$\delta(2)\backslash R$	32		48	32		48
Panel A. Wild cluster bootstrap						
1/2	0.06	$<^{***}$	0.47	0.06	$<^{***}$	0.35
	$\wedge^{***}$	$\angle$	$\wr$	$\wedge^{***}$	$\angle$	$\wr$
2/3	0.36	$<$	0.44	0.30	$<$	0.35
Panel B. Player-clustered probit						
1/2	0.06	$<^{***}$	0.47	0.06	$<^{***}$	0.35
	$\wedge^{***}$	$\angle$	$\wr$	$\wedge^{***}$	$\angle$	$\wr$
2/3	0.36	$<$	0.44	0.30	$<$	0.35
Panel C. Session-level rank-sum						
1/2	0.06	$<^{***}$	0.47	0.06	$<^{***}$	0.35
	$\wedge^{***}$	$\angle$	$\wr$	$\wedge^{***}$	$\angle$	$\wr$
2/3	0.36	$<$	0.44	0.30	$<$	0.35

Notes: Each panel reproduces Table 6 with significance from a different procedure. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table A2: Cooperation Rates in First Repeated Game — alternative inference

	First round			All rounds		
$\delta(2)\backslash R$	32		48	32		48
Panel A. Wild cluster bootstrap						
1/2	0.38	$<^{*}$	0.62	0.16	$<^{***}$	0.53
	$\wr$	$\angle^{**}$	$\wr$	$\wedge^{*}$	$\angle^{**}$	$\wr$
2/3	0.38	$<^{***}$	0.65	0.31	$<^{**}$	0.49
Panel B. Player-clustered probit						
1/2	0.38	$<^{**}$	0.62	0.16	$<^{***}$	0.53
	$\wr$	$\angle^{**}$	$\wr$	$\wedge^{***}$	$\angle^{**}$	$\wr$
2/3	0.38	$<^{***}$	0.65	0.31	$<^{**}$	0.49
Panel C. Session-level rank-sum						
1/2	0.38	$<$	0.62	0.16	$<^{**}$	0.53
	$\wr$	$\angle^{**}$	$\wr$	$\wedge$	$\angle^{**}$	$\wr$
2/3	0.38	$<^{**}$	0.65	0.31	$<$	0.49

Notes: Each panel reproduces Table 7 with significance from a different procedure. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

B Risk Aversion

I measure individual risk attitudes using the Bomb Risk Elicitation Task (BRET) of Crosetto and Filippin (2013) with oTree implementation of Holzmeister and Pfurtscheller (2016), administered after the prisoner's dilemma game. In the BRET, subjects sequentially collect boxes from a grid in which one box conceals a hidden bomb. Collecting additional boxes raises potential earnings but also increases the probability of losing all accumulated payoff. The number of boxes collected thus provides a continuous, incentive-compatible measure of individual risk tolerance.

Table B1 reports distributional statistics for the number of boxes collected across the four experimental treatments. The distributions appear broadly similar.

Table B1: Summary Statistics of Boxes Collected, by Treatment

	Mean	SD	p25	Median	p75	Min	Max	N
$\delta_2 = 2/3,\ R = 48$	30.00	17.59	14.00	29.00	41.50	4	66	52
$\delta_2 = 2/3,\ R = 32$	34.41	17.09	23.50	33.50	49.00	1	70	56
$\delta_2 = 1/2,\ R = 48$	31.52	20.59	13.50	29.50	48.50	2	93	60
$\delta_2 = 1/2,\ R = 32$	36.66	18.05	25.00	35.00	49.00	3	99	56
Total	33.17	18.49	20.00	33.00	47.50	1	99	224

Notes: The table reports distributional statistics for the number of boxes collected in the BRET across the four experimental treatments.

To assess whether individual risk attitudes predict initial cooperation decisions, Table B2 reports the estimated marginal effects of the BRET measure on the probability of cooperating in the first round of each match, estimated separately for each treatment. All specifications include match fixed effects. Standard errors are clustered at the individual level. Across all four treatment, the estimated marginal effect of boxes collected on first-round cooperation is small and statistically indistinguishable from zero.

Table B2: Marginal Effects of Risk Attitudes on First-Round Cooperation

	$\delta(2) = 1/2$		$\delta(2) = 2/3$	
	$R = 32$	$R = 48$	$R = 32$	$R = 48$
Boxes collected	0.002	0.000	-0.002	0.001
	(0.001)	(0.003)	(0.004)	(0.003)
Observations	1952	1955	1560	1678

Notes: The table reports marginal effects from probit regressions of first-round cooperation on the number of boxes collected in the BRET. Each column corresponds to a separate treatment subsample. Match number fixed effects are included in all specifications. Standard errors are clustered at the individual level. Subjects collecting fewer than five boxes are excluded from the estimation sample; their exclusion does not affect the results.

Taken together, risk aversion therefore does not appear to confound the treatment comparisons reported in the main analysis.

C Strategies

C.1 AD

The Always Defect (AD) strategy is defined as:

$$s_d(h^t) = D, \quad \forall h^t.$$

To assess whether AD can be supported as a SPE, I again apply the one-shot deviation principle. Under AD, both players defect in every period, yielding a constant per-period payoff of P. The continuation probability remains at $\delta(0)$.

The expected payoff from following AD is:

$$U_i = P + \delta(0)P + \delta(0)^2 P + \cdots = \frac{P}{1-\delta(0)}.$$

Now consider a one-shot deviation in which a player unilaterally cooperates in a given round. In that case, the deviator receives the sucker's payoff S, while the opponent defects and the continuation probability rises to $\delta(1)$, since only one player cooperated. From the next period onward, both players revert to mutual defection, yielding continuation payoffs:

$$U_i' = S + \delta(1)\left(\frac{P}{1-\delta(0)}\right).$$

A deviation is not profitable if:

$$\frac{P}{1-\delta(0)} \geq S + \delta(1) \cdot \frac{P}{1-\delta(0)},$$

which simplifies to:

$$\frac{1-\delta(1)}{1-\delta(0)} \cdot P \geq S. \tag{AD}$$

Finally, it is not optimal to cooperate on any out-of-equilibrium history because the game follows the similar structure as described above.

Condition AD characterizes when AD is a sustainable SPE: the gain from unconditional defection must outweigh the short-term cost of increasing continuation probability and receiving the sucker's payoff from a cooperative deviation.

C.2 Grim

The Grim Trigger strategy is defined as follows:

$$s_g(h^t) = \begin{cases} C, & \text{if } a = (C, C)\ \forall a \in h^t \text{ or if } h^t = \varnothing \\ D, & \text{otherwise.} \end{cases}$$

To determine whether Grim constitutes a subgame perfect equilibrium (SPE) under the cooperation-contingent continuation structure, I apply the one-shot deviation principle.

Case 1: Deviation after mutual cooperation (C, C). The expected on-path payoff is:

$$U_i = R + \delta(2)R + \delta(2)^2 R + \delta(2)^3 R + \cdots = \frac{R}{1 - \delta(2)}.$$

Consider a one-stage deviation in which player i defects while the opponent cooperates. This yields the temptation payoff T in the current period, followed by a transition to mutual defection. Under Grim, defection triggers permanent punishment. Since the defection was unilateral, the continuation probability falls to $\delta(1)$, and subsequent payoffs are from mutual defection, i.e., P, with continuation probability $\delta(0)$. The expected payoff from deviating is:

$$U_i' = T + \delta(1)\left(P + \delta(0)P + \delta(0)^2 P + \dots\right) = T + \delta(1) \cdot \frac{P}{1 - \delta(0)}.$$

Thus, Grim is a SPE if no profitable one-shot deviation exists, i.e., if:

$$\frac{R}{1 - \delta(2)} \geq T + \delta(1) \cdot \frac{P}{1 - \delta(0)}. \tag{2}$$

Case 2: Deviation after (C, D), (D, C), or mutual defection (D, D). The expected on-path payoff following defection is:

$$U_i = P + \delta(0)P + \delta(0)^2 P + \delta(0)^3 P + \cdots = \frac{P}{1 - \delta(0)}.$$

Consider a one-stage deviation in which player i cooperates while the opponent defects. This yields the sucker payoff S in the current period, followed by a mutual defection. Since the cooperation was unilateral, the continuation probability jumps to $\delta(1)$, and subsequent payoffs are from mutual defection, i.e., P, with continuation probability $\delta(0)$. The expected payoff from deviating is:

$$U_i' = S + \delta(1)\left(P + \delta(0)P + \delta(0)^2 P + \dots\right) = S + \delta(1) \cdot \frac{P}{1 - \delta(0)}.$$

Thus, Grim is a SPE after (D,D) history if:

$$\frac{P}{1-\delta(0)} \geq S + \delta(1) \cdot \frac{P}{1-\delta(0)}$$

$$\frac{1-\delta(1)}{1-\delta(0)} \cdot P \geq S. \tag{3}$$

So, **Grim is a SPE if condition AD holds and**

$$\frac{R}{1-\delta(2)} \geq T + \delta(1) \cdot \frac{P}{1-\delta(0)}. \tag{Grim}$$

C.3 AC

I also consider the Always Cooperate (AC) strategy, motivated by the possibility that players may choose to cooperate in order to increase the probability of continuation, even when doing so yields a lower immediate payoff. In environments with cooperation-contingent continuation, mutual cooperation raises the likelihood of future interaction, potentially offsetting the short-term temptation to defect.

The AC strategy is defined as:

$$s_c(h^t) = C, \quad \forall h^t.$$

To evaluate whether AC can be supported as a subgame perfect equilibrium (SPE), I apply the one-shot deviation principle. If all players follow AC, the expected payoff from perpetual mutual cooperation is:

$$U_i = \frac{R}{1-\delta(2)}.$$

Now consider a one-shot deviation to defection, which yields a temptation payoff T in the current round. Since only one player defects, the continuation probability in the next round drops to $\delta(1)$, after which players resume mutual cooperation. The deviator's expected payoff is:

$$U_i' = T + \delta(1) \cdot \frac{R}{1-\delta(2)}.$$

The deviation is not profitable if:

$$\frac{R}{1-\delta(2)} \geq T + \delta(1) \cdot \frac{R}{1-\delta(2)},$$

which simplifies to:

$$\frac{1-\delta(1)}{1-\delta(2)} \cdot R \geq T. \tag{AC}$$

Finally, it is not optimal to cooperate on any out-of-equilibrium history because the game follows

the similar structure as described above.

Condition AC characterizes when the AC strategy is sustainable as a SPE. In such cases, the higher probability of continuation from mutual cooperation may outweigh the short-run gains from defection, thereby supporting unconditional cooperation in equilibrium.

C.4 Tit-for-Tat (TFT)

The Tit-for-Tat (TFT) strategy is defined as:

$$s_{\text{TFT}}(h^t) = \begin{cases} C, & \text{if } h^t = \varnothing \text{ or } a^{t-1} \in \{(C,C),(D,C)\}, \\ D, & \text{otherwise.} \end{cases}$$

When both players adhere to TFT, they cooperate in every period, sustaining the cooperative path (C, C). I examine two key deviations to assess whether TFT can be supported as a subgame-perfect equilibrium (SPE).

Case 1: Deviation after mutual cooperation (C, C). Suppose one player unilaterally deviates by defecting after a history ending in $a^{t-1} = (C, C)$, while the opponent continues playing TFT. This leads to a cycle of alternating play: $(D, C), (C, D), (D, C), \ldots$. The deviator receives T and S in alternating rounds.

The expected payoff from such a deviation is:

$$U_i' = T + \delta(1)S + \delta(1)^2 T + \delta(1)^3 S + \cdots = \frac{T + \delta(1)S}{1 - \delta(1)^2}.$$

The payoff from mutual cooperation under TFT is:

$$U_i = \frac{R}{1 - \delta(2)}.$$

Thus, for cooperation to be sustainable after the cooperative path, the following condition must hold:

$$\frac{R}{1 - \delta(2)} \geq \frac{T + \delta(1)S}{1 - \delta(1)^2} \tag{4}$$

Case 2: Deviation after unilateral cooperation by the opponent (D, C). Suppose a player defects in a round following $a^{t-1} = (D, C)$. Under TFT, the opponent now plays D, and the deviator plays D, so the play in this and following rounds is going to be (D, D) in the next round.

The expected payoff from following TFT after (D, C) is the same alternating stream:

$$U_i = \frac{S + \delta(1)T}{1 - \delta(1)^2}.$$

A deviation, which brings immediate mutual defection, yields:

$$U_i' = \frac{P}{1-\delta(0)}.$$

Therefore, for TFT to be incentive-compatible after this history, I require:

$$\frac{S+\delta(1)T}{1-\delta(1)^2} \geq \frac{P}{1-\delta(0)} \tag{5}$$

Case 3: Deviation after unilateral defection by the opponent (C,D)**.** Suppose a player cooperate in a round following $a^{t-1} = (C,D)$. Under TFT, the opponent now plays C, and the deviator plays C, so the play in this and following rounds is going to be (C,C) in the next round.

The expected payoff from following TFT after (C,D) is the alternating stream:

$$U_i = \frac{T+\delta(1)S}{1-\delta(1)^2}$$

A deviation, which brings immediate mutual defection, yields:

$$U_i' = \frac{R}{1-\delta(2)}$$

Therefore, for TFT to be incentive-compatible after this history, I require:

$$\frac{T+\delta(1)S}{1-\delta(1)^2} \geq \frac{R}{1-\delta(2)} \tag{6}$$

Case 4: Deviation after mutual defection (D,D)**.** Suppose a player cooperates in a round following $a^{t-1} = (C,D)$. Under TFT, the opponent now plays C, and the deviator plays D. This leads to a cycle of alternating play: $(C,D),(D,C),(C,D),\ldots$. The deviator receives S and T in alternating rounds.

The payoff from mutual defection under TFT is:

$$U_i = \frac{P}{1-\delta(0)}$$

A deviation, which brings alternation, yields:

$$U_i = \frac{S+\delta(1)T}{1-\delta(1)^2}$$

Therefore, for TFT to be incentive-compatible after this history, I require:

$$\frac{P}{1-\delta(0)} \geq \frac{S+\delta(1)T}{1-\delta(1)^2} \tag{7}$$

Equilibrium implication. Together, conditions imply that TFT can be supported as a subgame-perfect equilibrium if and only if:

$$\begin{cases} \dfrac{R}{1-\delta(2)} \geq \dfrac{T+\delta(1)S}{1-\delta(1)^2} \\ \dfrac{P}{1-\delta(0)} \leq \dfrac{S+\delta(1)T}{1-\delta(1)^2} \\ \dfrac{R}{1-\delta(2)} \leq \dfrac{T+\delta(1)S}{1-\delta(1)^2} \\ \dfrac{P}{1-\delta(0)} \geq \dfrac{S+\delta(1)T}{1-\delta(1)^2} \end{cases} \tag{8}$$

Which holds if and only if:

$$\begin{cases} \dfrac{R}{1-\delta(2)} = \dfrac{T+\delta(1)S}{1-\delta(1)^2} \\ \dfrac{P}{1-\delta(0)} = \dfrac{S+\delta(1)T}{1-\delta(1)^2} \end{cases} \tag{TFT}$$

D Comparing to Kochov and Song (2023)

Following Kochov and Song (2023), player i exhibits decreasing marginal impatience (DMI) if for every $a, a' \in A^2$ and corresponding $c = c(a), c' = c(a')$,

$$v_i(a, a, \dots) > v_i(a', a', \dots) \quad \Leftrightarrow \quad \delta(c) > \delta(c') \tag{DMI}$$

Vice versa, the player i exhibits increasing marginal impatience (IMI) if for every $a, a' \in A^2$ and corresponding $c = c(a), c' = c(a')$,

$$v_i(a, a, \dots) > v_i(a', a', \dots) \quad \Leftrightarrow \quad \delta(c) < \delta(c') \tag{IMI}$$

The results for Prisoners' Dilemma from Kochov and Song (2023) are not applied to the present case because of the following propositions:

Proposition 1. *The game does not follow DMI.*

Proof. The DMI condition states that it holds for every $a, a' \in A^2$; therefore, it suffices to provide a counterexample. Consider $a = (D, C)$, $a' = (C, C)$, and parameters such that AC is not an equilibrium strategy. Then $c = 1$ and $c' = 2$, and it can be calculated that $v_i(a, a, \dots) > v_i(a', a', \dots)$, but $\delta(c) < \delta(c')$, which contradicts the DMI condition. □

Proposition 2. *The game does not follow IMI.*

Proof. The IMI condition states that it holds for every $a, a' \in A^2$; therefore, it suffices to provide a counterexample. Consider $a = (C, C)$ and $a' = (C, D)$, then $c = 2$ and $c' = 1$. It can be calculated that $v_i(a, a, \dots) > v_i(a', a', \dots)$, but $\delta(c) > \delta(c')$, which contradicts the IMI condition. □

E Instructions for one of the treatments

You are about to participate in an experiment on decision-making. Your earnings will depend partly on your decisions, partly on the decisions of others, and partly on chance. Please turn off your cell phones and similar devices now.

The entire session will take place through computer terminals, and all interaction between participants will occur through the computers. Please do not talk or attempt to communicate with other participants during the session.

This experiment consists of two parts; these instructions apply to the first part. Once this part is over, you will receive instructions for the next part. Your decisions in this part will not influence the other part.

If you have any questions, you can raise your hand, and the experimenter will come to assist you.

General Instructions for Part 1

In this experiment, you will be asked to make decisions in several rounds. You will be randomly paired with another person for a sequence of rounds. Each sequence of rounds is referred to as a match, and you will interact with the same person for that entire match. After each match concludes, a new match will begin, and you will again be randomly paired with another participant.

The points you can obtain in each round of a match depend on your choice and the choice of the person you are paired with. Each round you and the other person will choose a number (1 or 2) by selecting the corresponding button on the screen. The table below represents all the possible outcomes:

		Other's choice	
		1	**2**
Your choice	**1**	48, 48	12, 50
	2	50, 12	25, 25

The table shows the points associated with each combination of your choice and the choice of the person you are paired with. The first entry in each cell represents the points you obtain for that round, while the second entry represents the points obtained by the person you are paired with.

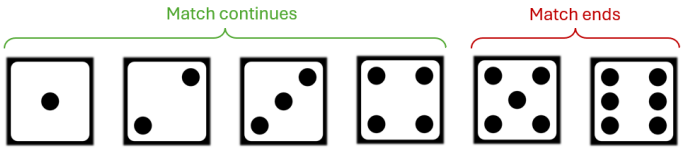
That is, in each round of a match, if:

- (1,1): Your choice is 1, and the other's choice is 1. You earn 48, and the other person earns 48.
- (1,2): Your choice is 1, and the other's choice is 2. You earn 12, while the other person earns 50.
- (2,1): Your choice is 2, and the other's choice is 1. You earn 50, while the other person earns 12.
- (2,2): Your choice is 2, and the other's choice is 2. You earn 25, and the other person earns 25

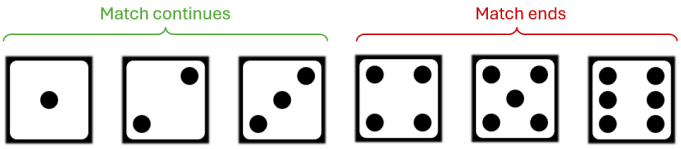
After each round, you will be reminded of your choice, and you will be shown the choice of the person you were paired with and all choices during previous periods for a given match.

At the end of each round, there is a probability that the match will continue for at least another round. The probability will depend on your choice and the choice of the person you are paired with. The computer will roll a 6-sided die in every round to decide whether the match will continue for at least another round:

- **If your choice in the previous round was 1 and the other person's choice was also 1**, the match will continue for an additional round if the die lands on a 1, 2, 3, or 4. The match will end if the die lands on a 5 or 6.



- **If at least one of you chose 2**, the match will continue for an additional round if the die lands on a 1, 2, or 3. The match will end if the die lands on a 4, 5, or 6.



However, you will play every match in blocks of 3 rounds. At the end of each block, you will learn if the match ended for all pairs in the room in the previous block of 3 rounds or not. If it has not, you will play another block of 3 rounds. If all matches have ended in this block, you will see in which round your match had actually ended. You will be informed of the die rolls generated by the computer for each round at the end of a match. The final round of the match will be the first round in which the die lands on

- 5 or 6, if your choice in the previous round was 1 and the other person's choice was also 1
- 4, 5, or 6, if at least one of you chose 2.

Remember, if you are in round 2, the probability that it is the last round is the same as it would be if you were in round 9.

Once a match ends, you will be randomly paired with someone for a new match. You will not be able to identify who you have interacted with in previous or future matches.

Each part of the experiment will generate points that count towards your final payoff. In this part, one match will be randomly selected to count towards your final payoff.

Points earned in this match will be converted to dollars at a rate of 15 cents per point or 3 dollars per 20 points. You will receive an additional $5 show-up fee for your participation. You will only be informed of your payoffs at the end of the experiment.

Quiz

1. **True**/False: The first entry in each cell represents your payoff, while the second entry represents the payoff of the person you are matched with.
2. True/**False**: You will be paid for all matches.
3. Open box: If you choose 1 and the other person chooses 2, your payoff in this round will be.
4. Multiple Choice: You have already played 3 rounds and are now in round 4. In this round, you chose 2 and your opponent chose 1. Under which outcomes of the die roll does the game continue?
   (a) 1,2,5,6
   (b) **1,2,3**
   (c) 3,4,5,6
   (d) 1,2,3,4

Summary

Before we start, let me remind you that:

- You will be randomly paired with another person for a match. You will interact with the same person for the entire match.
- After each round, the computer will roll a fair 6-sided die to determine whether the match will continue for at least another round
- You will play each match in blocks of 3 rounds; and will be informed of whether the match had ended during the block at the end of these rounds. You will NOT receive any payoff from rounds you have played within a block after the match had ended.
- Your choice and the choice of the person you are paired with will be shown to both of you at the end of each round.
- Your payoff in each round depends on these choices.
- After a match is finished, you will be randomly paired with someone for a new match.

If you have any questions, you can raise your hand, and the experimenter will come to assist you.

General Instructions for Part 2

On the screen, you see a field composed of 100 boxes, as shown below (the numbers on each box will not be visible):

1	2	3	4	5	6	7	8	9	10
11	12	13	14	15	16	17	18	19	20
21	22	23	24	25	26	27	28	29	30
31	32	33	34	35	36	37	38	39	40
41	42	43	44	45	46	47	48	49	50
51	52	53	54	55	56	57	58	59	60
61	62	63	64	65	66	67	68	69	70
71	72	73	74	75	76	77	78	79	80
81	82	83	84	85	86	87	88	89	90
91	92	93	94	95	96	97	98	99	100

- There is also a **"Start"** button. Once the button is clicked, the experiment begins. Every two seconds, a box will be collected, starting with Box 1 (top left) and ending with Box 100 (bottom right).
- You earn **1 point** for every box that is collected. Once collected, the box becomes marked, and your earnings are updated accordingly. At any moment, in the information box, you can see the number of boxes collected so far.
- However, these earnings are only potential, because behind one of the boxes a **bomb** is hidden that will destroy everything you have collected in this part of the experiment. You do not know the location of the bomb, and even if you collect it, you will not know until the end of the experiment. Your task is to decide when to stop collecting. You stop the process by hitting *"Stop"* at any time, and you will not be able to continue once you hit this button.
- When you hit *"Solve"*, the interface will show whether the boxes you collected contained a bomb.
- If none of the boxes you have collected contain the bomb when you hit *"Stop"*, you will receive the points you have accumulated. However, if you have collected the box with the bomb, you will earn **0**. Remember, you will not be told whether a box you have collected contains the bomb until after you hit the *"Solve"* button. Therefore, the earnings displayed on the screen are only potential, and you will receive them only if none of the boxes you collected contained the bomb.
- Points earned in this match will be converted to dollars at a rate of 15 cents per point or **3 dollars per 20 points**.
- Remember, decisions in the first part do not influence this part.
- Please note that there will be no trial period. You will take only one decision, and this will be payoff relevant. You will see now the main interface of the experiment. Take your time to observe it, before you click on *"Start"*.