

Beliefs in Repeated Games: An Experiment*

Masaki Aoyagi
Osaka University

Guillaume R. Fréchet
NYU

Sevgi Yuksel
UCSB

February 13, 2024

Abstract

This paper uses a laboratory experiment to study beliefs and their relationship to action and strategy choices in finitely and indefinitely repeated prisoners' dilemma games. We find subjects' elicited beliefs about the other player's action are generally accurate despite some systematic deviations, and anticipate the evolution of behavior differently between the finite and indefinite games. We also use the elicited beliefs over actions to recover beliefs over supergame strategies played by the other player. We find these beliefs over strategies correctly capture the different classes of strategies played in each game, vary substantially across subjects, and rationalize their strategies.

JEL classification: C72, C73, C92.

Keywords: repeated game, belief, strategy, elicitation, prisoner's dilemma.

*We would like to thank the editors, anonymous referees and numerous seminar participants for helpful comments. We also wish to thank Hazem Alshaikhmubarak, Sonal Barve, and Caroline Zhang for excellent research assistance. We are grateful for the financial support from National Science Foundation (grants SES-1558857, SES-1225779 and SES-1558857), the CV Starr Center and the Center for Experimental Social Science (CESS), the JSPS (grant numbers: 21653016, 24653048, 15K13006, 22330061, 15H03328 and 15H05728), the Joint Usage/Research Center at ISER, Osaka University, and the International Joint Research Promotion Program of Osaka University for research support. We are responsible for all errors. Aoyagi: aoyagi@iser.osaka-u.ac.jp; Fréchet: frechette@nyu.edu; Yuksel: sevgi.yuksel@ucsb.edu.

1 Introduction

Equilibrium analysis assumes players have correct beliefs about the strategies of other players and they best respond to these beliefs. These assumptions may be particularly demanding in repeated games, where (i) strategies can be very complex, (ii) there can be multiplicity of equilibria, and (iii) learning is made difficult by the large number of possible histories. This paper uses a laboratory experiment to study the validity of these assumptions by constructing a novel data set that includes *beliefs* as well as *actions* in repeated prisoners' dilemma (PD) games. By making beliefs observable, our goal is to bring to light a key force at work in determining behavior in such games.

Our experiment on repeated PD games consists of two main treatments: the *Finite* game and the *Indefinite* game. Theory predicts the existence of a unique equilibrium with no cooperation in the Finite game, but the existence of a multitude of equilibria ranging from no cooperation to full cooperation for sufficiently patient players in the Indefinite game. This theoretical contrast between the two games provides a useful backdrop for the study of beliefs and their relationship to cooperation. Based on the literature, we select parameters so that these two games generate similar and high levels of round-one cooperation in the laboratory. The two treatments hence allow us to compare beliefs among subjects taking the same action in the same round potentially along the same history, and examine whether their strategic reasoning is similar or different across the two environments.

In a first foray into beliefs in repeated PD games, many questions are of interest. However, given the challenges associated with implementing both repeated games and eliciting beliefs in the laboratory, we have opted for simplicity whenever possible: we use games with perfect monitoring where the past actions of both players are perfectly observable, and only elicit (first-order) beliefs about stage actions.

Our analysis is on both beliefs about stage actions, or *round beliefs*, and beliefs about supergame strategies, or *supergame beliefs*, which are recovered from round beliefs using a novel method. Round beliefs are informative since they are a cross section of supergame beliefs and a more primitive record of subjects' strategic thinking in reaction to history of play and other features of the game. When compared with eliciting supergame beliefs, eliciting round beliefs is cognitively less demanding and requires no assumption about the underlying supergame strategies. It is also less likely to alter how subjects approach the strategic interaction. The method we use to recover supergame beliefs from round beliefs is in two steps: First, we type subjects by assigning them to one of the supergame strategies in a predefined consideration

set based on their stage actions. Second, we use elicited round beliefs to estimate, for each type separately, the supergame beliefs over strategies in this set.

We identify three classes of key results. First, beliefs are, broadly speaking, *accurate*. This is noticeable at many levels: Round by round, unconditional average round beliefs are close to empirical action frequencies. Round beliefs are also *history-dependent*. In round two, subjects display large changes in their beliefs that closely reflect the actual change in action frequencies. For instance, in both treatments, subjects who cooperate while their opponent defects in round one decrease their belief on the likelihood that their opponent cooperates by an average of more than 40 percentage points. Round beliefs are also *forward-looking*. Beliefs towards the end of the Finite game correctly anticipate that cooperation is substantially less likely, a pattern not displayed in the Indefinite game. The most striking example of this is that subjects in pairs that have jointly cooperated for seven rounds estimate the probability that the other will cooperate in round eight to be below 60% in the Finite game, but above 95% in the Indefinite game. Round beliefs are informed by past experience, but cannot be reduced to it. For instance, in more than three quarters of cases, subjects' round one beliefs differ from the cooperation rate they have experienced in earlier supergames by more than 10 percentage points. In fact, in 58 percent of cases beliefs are not even within plus or minus 20 percentage points of the cooperation rates experienced in earlier supergames.¹ As for supergame beliefs, they correctly anticipate the types of strategies prevalent in each environment. Specifically, their support includes conditionally cooperative strategies in both environments. These strategies are stationary in the Indefinite game but are non-stationary and switch to defection in the last few rounds in the Finite game.

Second, despite the aforementioned general accuracy, beliefs also display small but systematic deviations. Round beliefs are too optimistic towards the end of the Finite game and too pessimistic at the beginning of the Indefinite game. While, as mentioned above, supergame beliefs correctly capture the prevalent strategies in both environments, such beliefs are not necessarily perfectly calibrated to the actual frequency of strategies in the population. In the Finite game, this plays a key role in preventing complete unravelling of cooperation. In the Indefinite game, as particularly visible in our additional treatments where the stage game payoffs are less conducive to cooperation, this provides an explanation for why payoffs are inside the efficiency frontier.

¹This is consistent with the finding of Nyarko and Schotter [2002] who report that beliefs are not the average of past observations, or more precisely the γ -weighted empirical average [Cheung and Friedman, 1997].

Third, beliefs are remarkably heterogeneous across subjects. This heterogeneity is directly visible in the distribution of round one beliefs, but is also present in the supergame beliefs. Specifically, the supergame beliefs vary substantially across types, and this variation helps rationalize their strategy choice: for most types, their strategy is a best response (or close to a best response) to their supergame beliefs among the strategies in our consideration set.

How do these findings inform our understanding of behavior in these games? In the Finite game, we observe high cooperation and partial unravelling, behavior not predicted by theory. Much of this can be explained by beliefs that are just slightly over-optimistic about how much others will cooperate in the last few rounds of the game. In a game such as the one studied here (and in many similar games that have been studied before), even a little over-optimism can substantially weaken incentives to defect earlier, hindering unravelling of cooperation. Our estimates suggest that, for 80% of subjects, best responding to their beliefs translates into *cooperating too much* relative to the best response against the actual strategy distribution. In the Indefinite game, on the other hand, we observe that a variety of strategies persist in the long run. Our results connect heterogeneity in strategies to heterogeneity in beliefs, which in turn rationalize such strategies. To organize these observations, we provide a stylized model that borrows elements from the level- k models [Stahl and Wilson, 1994, Nagel, 1995, Stahl and Wilson, 1995, Camerer et al., 2004, Costa-Gomes et al., 2001] as well as the *gang of four* model [Kreps et al., 1982]. Our model illustrates how the key systematic deviations—over-optimism in later rounds of the Finite game and over-pessimism in initial rounds of the Indefinite game—can result from a common mistake where players believe others to be less strategically sophisticated than themselves.

The present paper contributes to a few strands of the literature. First, it contributes to the literature that studies consistency of beliefs and strategies. To the extent that this consistency has been studied experimentally, the focus has been on one-shot games, on which there are mixed results. For example, Nyarko and Schotter [2002] find that subjects, for the most part, best respond to their beliefs. Other papers, most notably Costa-Gomes and Weizsäcker [2008], focusing on different stage games, do not find such behavior to be as prevalent. Rey-Biel [2009], who reports a high fraction of best-response behavior, suggests that a general conclusion on consistency is difficult as it may depend on various features of the game. His results, however, indicate that best-response behavior may be higher in simple games. Review of the broader literature in Online Appendix A suggests an interesting pattern: Lower rates of best response are reported when the game is not played

multiple times or played with no feedback [Costa-Gomes and Weizsäcker, 2008, Hyndman et al., 2022, Danz et al., 2012, Rey-Biel, 2009]. Conversely, Hyndman et al. [2012a] show best response behavior to increase with experience in an experiment with feedback. Our contribution to this literature is to study consistency of beliefs and actions in the finitely and indefinitely repeated PD, which, as discussed earlier, poses unique challenges.² Despite these challenges, consistent with earlier results from experiments with repetition and feedback, we find behavior to be close to best response for a majority of our subjects.

Our results also speak to the rapidly growing literature on experiments with belief elicitation. Most of the papers in that literature examine beliefs in individual decision making settings (see Danz et al. [2020] for a recent review), and those that study beliefs in games mostly use one-shot games.³ Most closely related to the present paper are the experiments that elicit beliefs in the voluntary contribution mechanisms (VCM), which are social dilemmas [Gächter and Renner, 2010, Neugebauer et al., 2009, Fischbacher and Gächter, 2010]. Although some of the studies on the subject involve designs with fixed pairing and feedback as in the present experiment, to the extent that prior experiments on beliefs have induced repeated games in the laboratory, they do so assuming that incentives in static interactions remain unchanged in repeated play. As such, these papers do not consider supergame strategies based on dynamic incentives, or the possibility of learning over multiple supergames. We contribute to this literature by studying beliefs in repeated games while highlighting clearly important dynamic incentives.

We contribute to the experimental as well as theoretical literature on repeated games in a few different ways. First, in a closely related paper, Gill and Rosokha [2020] study indefinitely repeated (but not finitely repeated) PD games. Their subjects directly choose one alternative from a list of ten supergame strategies. By eliciting subjects' (supergame) beliefs over how the other player chooses from the same list in the first and last supergames, Gill and Rosokha [2020] study how supergame beliefs change with experience and personality traits. Their results show that beliefs respond to experience and are more accurate in the last supergame than in the first. Duffy et al. [2021] find that their subjects fail to best respond against robot players, which are known to follow the Grim trigger strategy in indefinitely repeated PD games. Beside studying behavior in two distinct types of repeated inter-

²The large number of histories can make learning difficult in repeated games. For example, in their final supergame, at least one third of subjects experience a history that is new to them.

³One exception is Davis et al. [2016] who elicit a crude measure of beliefs by asking subjects to guess the action of their opponent in an indefinitely repeated PD. Data (analyzed in their appendix) show a correlation between guesses and actions.

actions, our point of departure from those papers is to study repeated games without restrictions on behavior. Finally, also note that in Kreps et al. [1982], cooperation in the finitely repeated PD games is supported by the presence of an irrational type. Our finding that many subject types cooperate in the Finite game while attaching positive belief weight to conditionally cooperative strategies lends empirical support to the formulation of Kreps et al. [1982].

2 Strategies and Beliefs

The stage game is the standard prisoners' dilemma with two actions, C (cooperation) and D (defection). Let $A_i = \{C, D\}$ be the set of (stage) actions, and let $A = A_1 \times A_2$ be the set of action profiles with a generic element a . The stage-game payoffs $g_i(a)$ are given in Table 1. The horizon of the supergame (repeated game) is either finite or infinite. For $t = 1, 2, \dots$, history h^t of length t is a sequence of action profiles in rounds $1, \dots, t$. Let $H^t = A^t$ be the set of t -length histories. A player's (behavioral) *strategy* $\sigma_i = (\sigma_i^1, \sigma_i^2, \dots)$ is a mapping from the set of all possible histories to actions. $\sigma_i^1(a_i) \in [0, 1]$ denotes the probability of action a_i in round 1, and for $t \geq 2$ and history h^{t-1} , $\sigma_i^t(h^{t-1})(a_i) \in [0, 1]$ denotes the probability of action a_i in round t given history h^{t-1} . Let Σ_i denote the set of strategies of player i . In the supergame with finite horizon $T < \infty$, player i 's payoff under the strategy profile is the simple average of stage payoffs:

$$u_i(\sigma) = T^{-1} \sum_{t=1}^T E_\sigma [g_i(a^t)],$$

where E_σ is the expectation with respect to the probability distribution of $h^T = (a^1, \dots, a^T)$ induced by σ . In the supergame with infinite horizon, the players have the common discount factor $\delta < 1$, and their payoff is the average discounted sum of stage-game payoffs:

$$u_i(\sigma) = (1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} E_\sigma [g_i(a^t)].$$

We postulate that each subject i is endowed with a supergame strategy $\sigma_i \in \Sigma_i$ and a *subjective belief* about the supergame strategy played by the other player.⁴ Specifically, we suppose player i believes j 's strategy is randomly chosen from some

⁴This formulation follows Kalai and Lehrer [1993]. See also Nachbar [2005] for a similar framework to model a player's ability to best respond to his belief that is asymptotically correct.

finite subset Z_j of Σ_j according to a probability distribution \tilde{p}_i , which is referred to as player i 's (prior) *supergame belief*.⁵ One interpretation of \tilde{p}_i is that it represents i 's prior belief over the proportion of different strategies played by the other subjects.⁶

Note \tilde{p}_i can be updated after each round of play conditional on realized history of play. For each $t \geq 2$ and $h^{t-1} \in H^{t-1}$, we denote by $\tilde{p}_i^t = \tilde{p}_i(\cdot \mid h^{t-1})$ player i 's updated supergame belief about j 's strategy in round t given h^{t-1} . Associated with this is player i 's *round t belief* $\mu_i^t(h^{t-1})$, which describes his belief about j 's stage action in round t . More specifically, $\mu_i^t(h^{t-1})$ is the probability that i assigns to j 's choice of action C given h^{t-1} , and is related to \tilde{p}_i^t through

$$\mu_i^t(h^{t-1}) = \sum_{\sigma_j \in Z_j} \tilde{p}_i^t(\sigma_j) \sigma_j^{t-1}(h^{t-1})(C).$$

The belief-elicitation task in this experiment involves beliefs over stage actions. That is, the design elicits from each subject i , in each round t (conditional on history of play), his belief $\mu_i^t \equiv \mu_i^t(h^{t-1})$. For simplicity, we often refer to μ_i^t as a ‘‘belief.’’ In section 5, we recover the subjects' supergame beliefs \tilde{p}_i from the sequence of their elicited beliefs μ_i^1, μ_i^2, \dots

Player i 's *type* refers to his supergame strategy σ_i . In our estimation of supergame beliefs, we assume player i is *Bayesian* in the sense that his supergame belief $\tilde{p}_i(\cdot \mid h^{t-1})$ is updated according to Bayes rule after each history: for any $t \geq 1$ and $h^t = (h^{t-1}, a^t)$,

$$\tilde{p}_i^t(\sigma_j) = \frac{\tilde{p}_i^{t-1}(\sigma_j) \sigma_j^{t-1}(h^{t-1})(a_j^t)}{\sum_{\tilde{\sigma}_j \in Z_j} \tilde{p}_i^{t-1}(\tilde{\sigma}_j) \tilde{\sigma}_j^{t-1}(h^{t-1})(a_j^t)},$$

where beliefs in the first round are $\tilde{p}_i^1 = \tilde{p}_i$. Player i is *subjectively rational* if his supergame strategy σ_i best responds to his supergame belief \tilde{p}_i :

$$\sigma_i \in \operatorname{argmax}_{\tilde{\sigma}_i \in Z_i} \sum_{\sigma_j \in Z_j} \tilde{p}_i(\sigma_j) u_i(\tilde{\sigma}_i, \sigma_j).$$

Some of the key supergame strategies in our analysis are as follows. AC and AD are the strategies that choose C and D , respectively, for every history. σ_i is Grim if $\sigma_i^t(h^{t-1})(C) = 1$ if $h^{t-1} = ((C, C), \dots, (C, C))$ and $\sigma_i^t(h^{t-1})(C) = 0$ otherwise. σ_i is TFT if $\sigma_i^1(C) = 1$ and $\sigma_i^t(h^{t-1})(a_j^{t-1}) = 1$ for every h^{t-1} and $t \geq 2$. For $k = 1, 2, \dots$, σ_i is Tk, a *threshold strategy* with threshold k , if σ_i follows Grim for all $t < k$, and then switches to AD after round k .

⁵We use \tilde{p} instead of p to denote beliefs. In later sections, we use p to denote the actual distribution of strategies in the population.

⁶With random matching, i 's belief about the strategy played by his opponent in each supergame is equal to his belief about the proportion of strategies in the population.

3 Design

The experiment involves two main (between-subjects) treatments, which we refer to as the *Finite* and *Indefinite* games. Three important considerations (besides the aforementioned aim for simplicity) guides our experimental design.

1. *Comparing the Finite and Indefinite games and selecting parameters such that initial cooperation rates are high in both.* Papers such as Dal Bó [2005] shows that for many parameter combinations, in line with theory, initial cooperation is lower in Finite games than in Indefinite games. However, as reported in Embrey et al. [2018], there are also parameter combinations for which high cooperation is observed in the Finite game (see also Lugovskyy et al. [2020]). This is not in line with theory and more surprising. For that reason, using past experiments as guidance, we selected parameters that were expected to generate high round one cooperation for both Finite and Indefinite games. This allows us to study whether cooperation is driven by similar considerations across these two games. Furthermore, robustness treatments introduced in Section 7 allow for further study of the impact of changing parameters within a game type.

2. *Introducing belief elicitation while mitigating its impact on the subject’s play.* One very important concern is that asking for beliefs from the onset of the experiment may alter how subjects approach the strategic interaction. To minimize this possibility, we separate the experiment into two parts. First, subjects are presented with “standard” repeated PD experimental instructions that do not mention beliefs. Second, after four supergames, the experiment is paused, and instructions explaining the belief-elicitation procedures are given. This two-part approach draws on Dal Bó and Fréchette [2019] and Romero and Rosokha [2023], who do this for strategy elicitation.⁷ Although this means not having beliefs in the first supergames, in our opinion, introducing belief elicitation without impacting play is a key concern and warrants such caution. Importantly, our method of delaying belief elicitation to later supergames does seem to be successful in not impacting behavior (see Section 4.1).

3. *Allowing subjects to gain ample experience.* Prior research, both with finitely and indefinitely repeated PD games, show the importance of experience [Embrey et al.,

⁷ Dal Bó and Fréchette [2019] find that choices in their experiments with strategy elicitation (introduced after a period of play of the repeated PD) are similar to those without strategy elicitation. Romero and Rosokha [2023] also find that choices are unaffected with such a design. Experiments that immediately introduce strategy elicitation have reported different results. See for instance the 2016 working paper version of Romero and Rosokha [2018]. In that early implementation of strategy elicitation where the elicitation started from the beginning, they report lower cooperation rates when doing elicitation as compared to direct choice.

Table 1: Stage Game

In ECU		Normalized			
	C	D		C	D
C	51, 51	22, 63	C	1, 1	-1.41 $\bar{6}$, 2
D	63, 22	39, 39	D	2, -1.41 $\bar{6}$	0, 0

2018, Dal Bó and Fréchette, 2018], in that behavior evolves in important ways and subjects need time to understand dynamic incentives. This is why the focus of this paper is on experienced behavior (beliefs and actions) as observed towards the end of the sessions. This desire to have subjects play as many supergames as possible is one of the factors that increase the need for simplicity.

We now turn to the specifics of the experimental design. The left panel of Table 1 shows the stage game used in the experiment (in experimental currency units), whereas the right panel shows its normalized version.⁸ Instructions use neutral language. In the paper, we use *supergame* to refer to each repeated game played between two matched players, and *round* to refer to each play of the stage game. In the Finite game, each supergame ends after eight rounds, $T = 8$. In the Indefinite game, there is a $\frac{7}{8}$ probability after each round that the supergame will continue for an additional round, inducing an expected supergame length of eight rounds [Roth and Murnighan, 1978]. To ensure the observation of at least eight rounds of play, the indefinite treatment uses the *block random design* that lets subjects play for eight rounds for sure, and then informs them of if and when the supergame actually ended; if it has not ended, they subsequently make choices one round at a time.⁹

At the conclusion of each supergame, subjects are randomly re-matched to play a new supergame. After four supergames are played, subjects are given new instructions on the belief-elicitation task. This is the first time beliefs are mentioned to the subjects. From that point onward, each subject i is asked in every round t to state their round t belief μ_i^t as an integer between 0 and 100. The task is incentivized via the *binarized scoring rule*, which determines the likelihood that a subject wins 50 experimental currency units based on their response in this task and the realized

⁸The normalization facilitates comparison with prior studies. With normalization, we set the mutual cooperation payoff equal to 1 and the mutual defection payoff equal to 0. The normalized temptation payoff is hence $2 = (63 - 39)/(51 - 39)$ and the normalized sucker payoff is $-1.41 = (22 - 39)/(51 - 39)$.

⁹This method was first introduced in Fréchette and Yuksel [2017]. As in Vespa and Wilson [2019], we only use the method for the first *block*.

action choice of the matched subject.¹⁰ The belief question is presented on a separate screen after subjects have made their action decision for that round and before feedback is provided. We opted for this ordering to minimize the risk that the belief questions influence the way subjects play these games. This process continues until the first supergame to terminate after at least one hour of play has elapsed.

Although prior research on indefinite PDs has not found that risk aversion is an important determinant of choices [Dal Bó and Fréchette, 2018], risk preferences could, in principle, mediate the relation between beliefs and choices. For this reason, we also elicited subjects’ risk preferences at the end of each session using the bomb task [Crosetto and Filippin, 2013]. Instructions for this task were distributed after the completion of the last supergame. For the main treatments, we conducted eight sessions per treatment. The relatively large number of sessions per treatment is required for the estimation of beliefs over strategies as presented in Section 5. The supergames for the part with belief elicitation are separated into *early* and *late* (see Table 5 in Online Appendix B for more information). We use this categorization in the presentation of results, with most of the data analysis focusing on late supergames. We randomly chose one supergame without belief elicitation and one supergame with elicitation for payment, and paid subjects for the outcomes of all game rounds for those two supergames. We also paid subjects for the belief-elicitation task in one randomly selected round of one randomly selected supergame.¹¹

4 Results

4.1 Actions

For any supergame, denote by x_i^t the indicator of subject i ’s choice of C in round t , and by \bar{x}^t , the round t cooperation rate averaged over subjects. As will be clear

¹⁰Incentive compatibility of the binarized scoring rule is independent of a subject’s risk attitude (Allen [1987], McKelvey and Page [1990], Schlag and van der Weele [2013], and Hossain and Okui [2013]). We use the implementation outlined in Wilson and Vespa [2018].

¹¹To address hedging concerns, we chose the supergame for the belief-elicitation task from the supergames not used for the action task. In addition, as is typical in experiments eliciting beliefs, the rewards for the beliefs (either 0 and 50) are smaller than those for choices (between 176 and 505 for an eight round supergame). ECUs were translated into dollars at an exchange rate of 3 cents per point. Maximal ECU earnings from the bomb task were 99. All subjects also received a show-up fee of \$8. Earnings from the experiment varied from \$22.00 to \$63.75 (with an average of \$35.30). All instructions (available in Online Appendix F) were read aloud. The computer interface was implemented using zTree [Fischbacher, 2007] and subjects were recruited from UCSB students using the ORSEE software [Greiner, 2015].




Figure 1: Cooperation Rate over Supergames

from the context, the analysis in what follows sometimes aggregates \bar{x}^t over multiple supergames.

Figure 1 shows cooperation rates by supergame. Starting with the Finite game (the left panel), we observe relatively high initial (round one) cooperation rates slightly above 80%. Focusing on rounds > 2 , and dividing the sample into two cases, x_i^t following the other player's cooperation $a_j^{t-1} = C$ and those following other's defection $a_j^{t-1} = D$, we observe high cooperation rates following cooperation and low cooperation rates following defection. We also observe that the difference between those two averages, referred to as *responsiveness*, increases with experience. The cooperation rate in round eight is decreasing with experience and is low by the end (below 20%).


The right panel of Figure 1 presents the same statistics for the Indefinite game. In this case, and as with the Finite game, round-one cooperation rates are high (start slightly below 80% and increase to slightly above 80%). Cooperation rates following cooperation by the other are high, whereas cooperation rates following defection are low. Again, responsiveness increases with experience. However, in contrast to the Finite game, cooperation rates in round eight are high and increasing with experience.

In Online Appendix B, we provide further analysis and confirm that behavior along key dimensions in our experiment is qualitatively consistent with prior findings on these two games without belief elicitation. In summary, consistent with prior experiments with comparable parameters, the design successfully generates similar and high levels of round-one cooperation in both games. Also in line with prior findings, subjects display responsiveness that increases with experience. Furthermore, cooperation collapses at the end of the Finite game but persists in the Indefinite game. Finally, it is worth noting that when regressing round one cooperation on potentially relevant regressors, a dummy variable that takes value one when beliefs are elicited and zero otherwise is not statistically significant (see Table 6).

Result 1 We reproduce qualitative data patterns observed in previous experiments on Finite and Indefinite PD games, and find no indication of actions being impacted by belief-elicitation. In particular, our results confirm cooperation is history dependent in both games. Furthermore, cooperation evolves differently in both games: it collapses at the end only in the Finite game.

4.2 Consistency of Actions and Beliefs

Let $\bar{\mu}^t = \sum_{i=1}^n \mu_i^t$ denote the average of round t beliefs, which is aggregated over multiple supergames and/or over particular histories in what follows.



Late supergames.

Figure 2: Choices and Beliefs by Round


Putting beliefs and actions together reveals beliefs—on average—to be remarkably accurate, often tracking cooperation rates within a range of a few percentage points. Figure 2 shows for late supergames that the point estimate for average belief $\bar{\mu}^t$ is close to that for the average cooperation rate \bar{x}^t in each round t and that their confidence intervals display substantial overlap. When aggregated over all first eight rounds, the differences between action frequencies and beliefs are small, at less than one percentage point for Finite and two percentage points for the Indefinite game. This difference is not statistically different from 0 for the Finite game, but it is for the Indefinite game (although small in magnitude).^{12,13}

However, looking at each round separately, in both games we see a statistical difference between action frequencies and beliefs for rounds one through three. The difference is about four percentage points for each of the three rounds of the Finite game, whereas it is 11, 5.8, and 2.0 percentage points for the same rounds of the Indefinite game. In rounds seven and eight, we also see statistically significant differences between action frequencies and average beliefs for the Finite game. The difference is 9.5 and 3.1 percentage points for rounds seven and eight, respectively. In other rounds (rounds 4-6 of the Finite game and rounds 4-8 of the Indefinite game), beliefs and cooperation rates are not statistically different at the 10% level. In summary, to the extent that action frequencies and beliefs differ, the deviations are most prominent for late rounds in the Finite game and for early rounds in the Indefinite game.

So far in Figure 2, we considered only unconditional beliefs, but what about the subjects' ability to anticipate actions following specific histories? To consider histories with a sufficient number of observations, we examine this for round two. Figures 3 and 4 present the relevant data conditional on round-one histories (labeled with one's own action first followed by the opponent's action). In both games, we observe that beliefs quickly adjust in response to the other's action. In all cases, beliefs move in the correct direction from round one to round two. Furthermore, subjects are capable of very large adjustments in beliefs, sometimes of more than 50


¹²We perform the test on the difference between the opponent's action (coded as 1 for cooperate and 0 for defect) and the reported belief. Results are robust to including all observation rounds or only the first eight rounds.

¹³Throughout, when statistically significant is used without a qualifier, it refers to the 10% level. Here and elsewhere, unless noted otherwise, statistical tests involve subject-level random effects and session-level clustering (see Fréchet [2012] and Online Appendix A.4. of Embrey et al. [2018] for a discussion of issues related to hypothesis testing for experimental data). In the case of beliefs, as here, we use a tobit specification allowing for truncation. For tests of cooperation, we use a probit specification.



Own choice listed first.
 Late supergames.
 Percentage of cases under the title.
 By construction the 'Actions' rate is 1 in round 1 of the left panels and 0 in the right panels.

Figure 3: Beliefs Conditional on Round One Action Pair, Finite Games



Own choice listed first.
 Late supergames.
 Percentage of cases under the title.
 By construction the 'Actions' rate is 1 in round 1 of the left panels and 0 in the right panels.

Figure 4: Beliefs Conditional on Round One Action Pair, Indefinite Games

percentage points. Note that this provides clear evidence of subjects updating their beliefs about the future cooperativeness of their counterpart following a history in which defection is observed. It is also interesting that such beliefs become equally pessimistic regardless of which player has defected. Comparing the two figures, we see action frequencies and beliefs evolve in a similar fashion in all panels except for the top-left panel, which shows clear differences across the two treatments. In the Finite game, most of the initially cooperative interactions eventually break down, and this breakdown is mirrored by beliefs. In the Indefinite game, on the other hand, beliefs about cooperation are sustained if they survive the first round.

Our focus, so far, has been the accuracy of beliefs on average. We study the accuracy of beliefs at the individual level in Online Appendix B. Table 8 in this appendix reports the accuracy of beliefs at several different precision levels. For rounds one and two (separated by round one history), the table reports the share of subjects (i) who correctly identify cooperation as a likely, unlikely, or uncertain event (classification is based on whether both beliefs and average cooperation rate are higher than 66 percent, lower than 33 percent, or between 33 and 66 percent respectively); (ii) whose beliefs are within five or ten percentage points of the average cooperation rate. The majority of subjects have broadly accurate beliefs, namely their beliefs lie in the same tercile (likely/unlikely/uncertain) as the observed cooperation rate. Nonetheless, only a minority of subjects hold beliefs that are within 10 or 5 percentage points of the actual frequency: 14 and 7 percent (10 and 7 percent) respectively for round one in the Finite (Indefinite) game. The accuracy of beliefs increases substantially in round two, particularly after the most commonly observed histories.

As Figures 3 and 4 above show, supergames starting with joint cooperation are the most common. How do beliefs evolve on a mutual cooperation path? Figure 5 shows the average cooperation rates \bar{x}^t and average beliefs $\bar{\mu}^t$ along the history $h^{t-1} = ((C, C), \dots, (C, C))$. For example, a solid circle at round five indicates the empirical cooperation rate after four rounds of joint cooperation (close to 100% in both games). The most striking observation is the sharp decline in beliefs toward the end in the Finite game. That is, subjects (correctly) anticipate the increasing likelihood of defection from their opponent despite the fact that all choices up to that point were cooperative for both players.¹⁴ Nonetheless, we see clear evidence that subjects underestimate the degree to which cooperation drops from round 6 to 7: whereas beliefs are well calibrated in round 6 (within 1 percentage point of the empirical frequency), they show over-optimism (13 percentage points higher than the

¹⁴The decline in beliefs is not driven by selection: conditioning on subjects who remain on a cooperative path until the eighth round, beliefs decline from 89% in round 2 to 49% in round 8.




Figure 5: Cooperative Path (First Eight Rounds)

empirical frequency) in round 7, and become better calibrated in round 8 (within 4 percentage points). In summary, these findings suggest that although subjects anticipate the decline in cooperation, they underestimate the magnitude and foresee only 60% of the actual drop in cooperation. In the Indefinite game, on the other hand, beliefs and cooperation rates remain high as the supergame unfolds. We also note that these patterns are already visible in early supergames (see Figure 18 in Online Appendix B).

These observations suggest that the evolution of beliefs in the Finite game cannot simply be explained by heuristic models based on past actions (within a supergame). For example, if a subject always set his belief equal to his opponent’s action in the previous round, he would report beliefs for round 7 (in the Finite game) that are almost three times more over-optimistic than the ones we observe in the data. Clearly, beliefs in the Finite game change on a cooperative path with the length of the interaction, and hence are non-stationary.

Result 2 (1) Beliefs are accurate, on average, but show some systematic and persistent deviations: they are over-optimistic in later rounds of the Finite game and over-pessimistic in earlier rounds of the Indefinite game. (2) Beliefs respond to the history of play. (3) However, differences exist across games even along the same history. In particular, subjects correctly anticipate cooperation will break down despite a history of joint cooperation in the Finite game.

We now turn to the question of whether different actions are supported by different beliefs. We summarize key results here and refer the reader to Online Appendix B for detailed analysis. First, we study the degree to which round-one beliefs are predictive of round-one actions. Overall, subjects with optimistic beliefs are more likely to cooperate in both treatments, but round-one beliefs are more predictive of round-one actions in the Indefinite game than in the Finite game.¹⁵ Second, we study how the distribution of beliefs differs across these games in later rounds conditional on the subject’s action in that round. Once again, higher cooperation rates are associated with more optimistic beliefs in both games, but cooperation and defection in certain rounds are associated with different beliefs for Finite versus Indefinite games. In round eight, particularly, beliefs of the subjects who cooperate are statistically different across treatments ($p < 0.001$), as are those of the subjects who defect ($p = 0.065$). Specifically, subjects who defect in round eight of the Finite game are more pessimistic (on average) than those who do so in the Indefinite game. Furthermore, subjects who cooperate are more optimistic in the Indefinite game than those in the Finite game. In fact, even at the very beginning of the supergames, subjects who defect in the Indefinite game are more pessimistic about the probability that their opponent will cooperate than subjects in the Finite game ($p < 0.001$).¹⁶

Result 3 Beliefs correlate to actions, and more optimistic subjects are more likely to cooperate. The same-round belief can generate different actions in each game.

5 Beliefs over Supergame Strategies

The preceding section finds a link between actions and beliefs, but to study whether subjects’ behavior is a best response to their beliefs we need to move beyond beliefs over actions and instead consider beliefs over strategies. The goal of this section is to develop an estimation method that takes as an input beliefs over actions (the data

¹⁵The analysis of Online Appendix B includes kernel density estimates of the distribution of round-one beliefs μ_i^1 separated by treatment and by the subject’s own action a_i^1 in round-one (Figure 19) and regression analysis (Table 9) of the determinants of cooperation. These results also suggest that risk preferences have some limited predictive power for round-one choice in the Finite game (with the likelihood of cooperation decreasing with risk aversion). In randomly terminated PDs, Proto et al. [2019] does not find a significant effect of risk preference on the first choice of a session, but Proto et al. [2022] does.

¹⁶These patterns are clearly visible in Figure 20 (Online Appendix B), which plots the CDF of beliefs by action and treatment for each round. Table 10 in the same Appendix depicts the marginal impact of beliefs (and round number) on the likelihood of cooperation.

collected in the experiment) and translates this to beliefs over supergame strategies as an output. We then use this method to study how the strategy choice relates to beliefs. We do not make the claim that subjects reason in terms of strategies *per se*, but that we can potentially represent their behavior as such.

Our estimation strategy consists of two stages:

- (A) Classify subjects into types based on the actions they take.
- (B) Estimate beliefs over supergame strategies separately for each type.

It is important to highlight that the two stages use different data: typing is based on actions only and estimation of supergame strategies relies in beliefs (conditioning on history). Thus, the estimation method does not impose any structure between typing and estimation of beliefs over supergame strategies, allowing us to meaningfully study variation in beliefs over supergame strategies by type. Since our primary interest in this section is studying whether subjects' behavior can be rationalized as a best response to their beliefs, in stage (A) we will type subjects based on the strategy they are most likely to be playing. However, the belief estimation method described in (B), which is the main innovation of this section, can easily be paired with alternative typing procedures.

Here, we outline the intuition for the approach developed in (B) using a simplified example. Suppose we want to recover beliefs over strategies for one player (referred to as player 1) when the data available to us are round beliefs over actions elicited in one supergame (against player 2). For the purpose of the example, assume we know player 1 believes that player 2 uses one of only three strategies: AD, AC, or Grim. In round one, we observe player 1's unconditional belief that his opponent will start by cooperating: $\mu_1^1 = 0.6$. From this belief, we can already infer the probability player 1 associates with player 2 playing AD, because it starts by defection. That is, we can infer $\tilde{p}(AD) = 0.4$ and $\tilde{p}(AC) + \tilde{p}(Grim) = 0.6$. However, we cannot determine $\tilde{p}(AC)$ or $\tilde{p}(Grim)$ separately. To do so, we look at beliefs elicited in other rounds of the supergame. Assume that in round one, player 1 plays D and player 2 plays C. After observing this history, player 1 reports his round-two belief: $\mu_1^2 = 0.1$. Because player 2 started by playing C, player 1 now knows she is not playing AD. However, player 1's belief about whether player 2 will cooperate in round two can reveal information about whether he believes player 2's strategy is more likely to be AC or Grim. Note that after such a history of (D, C) , the two strategies indeed prescribe different actions: D for Grim and C for AC. Given μ_1^2 , we can recover (via Bayes' rule) that $\tilde{p}(AC) = 0.06$ and $\tilde{p}(Grim) = 0.54$. This method provides us with a roadmap for how we can recover ex-ante beliefs over strategies using data on beliefs

over stage actions elicited in each round of a supergame. In addition, we allow for players to believe others implement their strategies with error and that subjects may report their belief with some error.

The example above lays out the intuition behind our methodology as well as highlighting some of the challenges it presents. We outline below how we address these challenges.

- (1) Belief estimation in the example above relies on the assumption that the relevant strategies (over which subjects have beliefs) are known.¹⁷ How do we specify the relevant set of strategies for our data set? By now, a significant body of literature documents which strategies are used in repeated PD games. This literature guides how we construct the consideration set.
- (2) The example was constructed such that the data can easily separate the strategies considered; but in some cases, this can require specific histories that are not common and thus call for more data. To increase sample size, we pool data from multiple subjects. However, assuming all subjects share the same beliefs seems unreasonable. Instead, we group subjects according to the strategy that best describes how they play, referred to as their *type*. We assume subjects of the same type share the same beliefs.¹⁸

Below we describe the estimation strategy in detail before presenting results.

5.1 Typing of Subjects

We type subjects based on the strategy they are most likely to be playing. We do this in two steps: (1) We estimate the distribution of strategies used at the population

¹⁷Note that this is not a challenge unique to our study but one encountered by *any* study that presents analysis involving strategies or beliefs over strategies in repeated games. In the literature studying strategies in repeated games, such an assumption is introduced either at the design stage, for studies eliciting strategies directly, by restricting the set of strategies available to subjects (essentially reducing the repeated game to a simultaneous move game), or a similar assumption is made at the estimation stage by focusing the analysis on a set of strategies. One advantage of the later approach, as adopted in this paper, is that the data can always be reanalyzed under different assumptions on the set of strategies considered.

¹⁸To validate this assumption, we do the following exercise. We compute the spread of beliefs defined as the difference between the 25th and 75th percentiles of beliefs averaged over rounds and histories. We test whether the spread of beliefs is less among subjects that are of the same type relative to all others in the population. Out of the 10 types (to be defined later) observed in the Finite game and the eight types in the Indefinite game; only three of the 18 paired comparisons are not in line with the assumption that the spread in beliefs is less among subjects of the same type.

level; (2) For each subject, given their choices, we compute the posterior likelihood of playing each strategy using population-level estimates as a prior. We classify subjects into strategy types by identifying the highest posterior. See Section 7 for a discussion of alternative approaches.

Population-Level Estimates of Strategies

We first use the Strategy Frequency Estimation Method (SFEM) introduced in Dal Bó and Fréchette [2011] to estimate the distribution of strategies used. The method first specifies the set of candidate strategies and then estimates their frequencies in a finite-mixture model allowing for the possibility of implementation errors. We use a two-step procedure to determine the set of strategies in our analysis. This set consists of AD, AC, Grim, TFT, STFT, Grim2, and TF2T, as well as threshold strategies T8, T7, and T6.¹⁹ Online Appendix B describes the two-step procedure and defines all strategies considered in our analysis. Formally, the SFEM results provide two outputs p and β , both at the population level: p is a probability distribution over the set of strategies, and β is the probability that the choice corresponds to what the strategy prescribes. The method identifies the values of p and β that maximize the likelihood of the observed sequences of actions.

Classification of Subjects

We use the SFEM results to compute the Bayesian posterior that a subject is playing each of the candidate supergame strategies given the sequence of their actions. Each subject is associated with the supergame strategy that has the highest likelihood according to this posterior.²⁰

To demonstrate how this works, consider a simpler setup where the set Z of candidate strategies consists only of AD and AC. Assume the SFEM yields $p = (p_{AD}, p_{AC}) = (0.7, 0.3)$ and $\beta = 0.9$. The corresponding behavioral strategies (that allow for implementation errors) are then given by \widehat{AD} and \widehat{AC} .

Consider a subject who, over multiple supergames consisting of 24 rounds in total, cooperates in 20 rounds and defects in four rounds. Given p and β , we can calculate the Bayesian posterior that this subject is playing \widehat{AD} versus \widehat{AC} . In fact,

¹⁹We refer to Grim, TFT, Grim2, and TF2T as conditionally cooperative strategies: they begin with cooperation and switch to defection only after some histories involving defection.

²⁰This approach allows for comparison to the many previous papers that do strategy estimation using SFEM. For instance, Dvorak [2020] recently provides a R-package for easy implementation of SFEM using the EM algorithm, which includes a very similar typing procedure.

the posterior that the subject is playing \widehat{AD} is $\frac{p_{AD}\beta^4(1-\beta)^{20}}{p_{AD}\beta^4(1-\beta)^{20}+p_{AC}\beta^{20}(1-\beta)^4}$, which is close to 0, whereas the posterior that he is playing \widehat{AC} is close to 1. Consequently, this subject would be typed as playing AC.

5.2 Estimating Supergame Beliefs

For each type in our data, we estimate their supergame beliefs over strategies \tilde{p} , as well as parameters $\tilde{\beta}$ and ν .²¹ Specifically, \tilde{p} is a probability distribution over the set $\tilde{Z}^{\tilde{\beta}}$, which has one-to-one correspondence with the set Z of candidate strategies used in the SFEM as follows: for each $\sigma_j \in Z$, $\tilde{\sigma}_j \in \tilde{Z}^{\tilde{\beta}}$ is a stochastic version of σ_j in the sense that at each history, $\tilde{\sigma}_j$ chooses the same action as σ_j with probability $\tilde{\beta}$, but chooses the other action by error with probability $1 - \tilde{\beta}$.

Note that \tilde{p} and $\tilde{\beta}$ jointly pin down beliefs over stage actions given each history. For illustration, suppose again that the set Z of candidate strategies consists only of AD and AC so that $\tilde{Z}^{\tilde{\beta}}$ consists of their randomized versions \widehat{AD} and \widehat{AC} for $\tilde{\beta} = 0.9$. It then follows that the round-one belief μ_i^1 equals $\tilde{p}_{\widehat{AD}} \times 0.1 + \tilde{p}_{\widehat{AC}} \times 0.9$. If the subject observes $a_j^1 = C$ in the first round, by Bayes' rule, his belief in round two will increase to

$$\left(\frac{\tilde{p}_{\widehat{AD}} \times 0.1}{\tilde{p}_{\widehat{AD}} \times 0.1 + \tilde{p}_{\widehat{AC}} \times 0.9} \right) 0.1 + \left(\frac{\tilde{p}_{\widehat{AC}} \times 0.9}{\tilde{p}_{\widehat{AD}} \times 0.1 + \tilde{p}_{\widehat{AC}} \times 0.9} \right) 0.9.$$

The third parameter ν represents potential errors in the reporting of beliefs. Formally, if a subject's belief in any round t (implied by \tilde{p} and $\tilde{\beta}$) is μ_i^t , we assume his reported belief is distributed according to the logistic distribution with mean μ_i^t and variance ν truncated to the unit interval. For each type, we identify the values of \tilde{p} , $\tilde{\beta}$, and ν that maximize the likelihood of the sequence of elicited beliefs in all rounds of late supergames.

5.3 Results

Population-level Estimates of Strategies and Classification of Subjects

Table 12 in Online Appendix C presents the estimation results of the distribution of strategies at the population-level (in columns 2 and 5) sorted by prevalence. The

²¹The variables with tilde are estimates about beliefs and distinguished from the corresponding SFEM estimates of strategies.

results are consistent with prior evidence on strategy choice in repeated PD: Threshold strategies are important in the Finite game [Embrey et al., 2018], and AD, Grim, and TFT account for a majority of the strategies in the Indefinite game [Dal Bó and Fréchette, 2018].²²

In the Finite game, T7 and T8 account for a little over half of the strategies, and they, along with AD, make up two thirds of the choices. Another threshold strategy, T6, is also in the top 5 at 8%. Additionally, TFT and Grim are commonly used strategies (at the 4th and 6th positions).

In the Indefinite game, conditionally cooperative strategies dominate, with TFT and Grim representing more than half of the choices. The lenient versions of Grim and TFT are also among the popular strategies, accounting together for 21% of the choices. Together these four account for more than two thirds of the strategies. Other prominent strategies are AC and AD, two unconditional strategies, representing 20% of the choices. All other strategies are at most 4% each, and the threshold strategies are almost completely irrelevant. Together, conditionally cooperative strategies account for 75% of the data (by contrast, these strategies represent only 21% of the data in the Finite game).

Table 12 also reports in the third and sixth columns the complete results of the typing exercise. The type shares are largely similar to the population estimates from SFEM. However, we also observe some differences. In particular, in the Indefinite game, the fraction of subjects typed as TFT is greater than the fraction of TFT in the population.²³ Clearly, the smaller the fraction of subjects of a given type, the less reliable their belief estimates will be.

Result 4 We reproduce results about strategy choices observed in previous finitely and indefinitely repeated PD games. In particular, our results confirm strategic heterogeneity exists within and across treatments. In the Finite game, subjects mostly use threshold strategies, whereas in the Indefinite game, they mostly rely on conditionally cooperative strategies.

²²Online Appendix C also reports SFEM results for early supergames (the changes are presented in Figure 27). Consistent with Embrey et al. [2018] those results show that threshold strategies increase with experience in the Finite game.

²³Two potential sources for such differences are possible. First, and simply mechanically, some subjects play more supergames than others; the fraction of subjects corresponding to a type can differ from the population (over supergames) fraction of that strategy. Second, imagine a data set where a large fraction of subjects is estimated to play TFT, and a smaller fraction is estimated to play Grim. However, there are subjects whose actions are equally consistent with Grim and TFT. Our method would type those subjects according to the prior. See Section 7.1 for more on this.

Estimates of Beliefs over Strategies

A summary of these estimation results are reported in Tables 2 and 3, with the complete results provided in Online Appendix C. Note that some types are not observed frequently enough to allow for estimation, which is the case whenever only 1% of subjects are of a certain type. In addition, there is sometimes insufficient variation to separate the beliefs with respect to some of the strategies. In those cases, we set the least popular strategies (according to SFEM) to zero and “assign” the belief to the more popular strategy. This applies to only three of the 84 estimates reported in Tables 2 and 3. The rows are sorted by frequency of the strategy, and the columns are sorted by average belief (i.e., the first strategy for which we report beliefs is the one that, on average, subjects put the most weight on).

Table 2: Beliefs over Strategies in the Finite Game

Type	Share		Estimated Beliefs - \tilde{p}								ν	$\tilde{\beta}$
	SFEM	Typing	T7	T8	Grim	TFT	AD	TF2T	Grim2	Other		
T7	0.30	0.35	0.43	0.39	0.18	0.00	0.00	0.00	0.00	0.00	0.04	1.00
T8	0.22	0.20	0.00	0.50	0.04	0.01	0.09	0.15	0.21	0.00	0.04	1.00
AD	0.12	0.12	0.75	[0.00]	[0.00]	0.00	0.07	0.00	0.00	0.18	0.06	1.00
TFT	0.09	0.12	0.00	0.33	0.00	0.53	0.11	0.00	0.00	0.03	0.05	1.00
T6	0.08	0.08	0.99	0.00	[0.00]	0.00	0.00	0.00	0.00	0.00	0.03	1.00
Grim	0.08	0.02	0.00	0.22	0.17	0.16	0.34	0.01	0.00	0.10	0.07	1.00
Other	0.11	0.11	0.01	0.14	0.30	0.26	0.01	0.09	0.05	0.11		
All			0.30	0.29	0.11	0.09	0.06	0.05	0.05	0.04		

Estimation on late supergames out of 10 strategies: AD, AC, Grim, TFT, STFT, T8-T6, Grim2, and TF2T.
 Rows, top 6 *played* strategies. Columns, top 7 *believed* strategies.
 Estimates in [square brackets] are not estimated due to collinearity.
 SFEM estimate for β is 0.94. Complete results in Table 15.

Tables 2 and 3 reveal important differences in beliefs between the Finite and Indefinite games. The bottom row of each table presents (weighted) average beliefs over strategies. In the Finite game, subjects believe others are most likely to use threshold strategies (T7 and T8 account for 59%), whereas in the Indefinite game, they believe others are most likely to play conditionally cooperative strategies (Grim and TFT have together 46%). That is, at least in this respect, subjects’ beliefs are in line with behavior in both games: subjects *correctly* anticipate the most popular class of strategies to be different between the games (threshold vs. conditionally cooperative). Furthermore, looking at the first two rows of each table, and focussing on the two most common strategies, we see evidence of substantial heterogeneity in beliefs between types (in the same game). For instance, T8 types in the Finite game

Table 3: Beliefs over Strategies in the Indefinite Game

Type	Share		Estimated Beliefs - \tilde{p}								ν	$\tilde{\beta}$
	SFEM	Typing	Grim	TFT	TF2T	AC	AD	Grim2	STFT	Other		
TFT	0.36	0.59	0.28	0.25	0.19	0.00	0.08	0.14	0.05	0.00	0.01	1.00
Grim	0.18	0.09	0.80	0.13	0.02	0.00	0.00	0.05	0.00	0.00	0.06	1.00
Grim2	0.11	0.11	0.22	0.00	0.23	0.23	0.00	0.31	0.00	0.00	0.02	1.00
AC	0.11	0.05	0.00	0.20	0.00	0.80	0.00	0.00	0.00	0.00	0.10	1.00
TF2T	0.10	0.01	0.33	0.00	0.40	0.27	0.00	0.00	0.00	0.00	0.01	1.00
AD	0.09	0.10	0.00	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.04	1.00
Other	0.05	0.05	0.29	0.00	0.00	0.00	0.39	0.13	0.00	0.00		
All			0.31	0.14	0.14	0.14	0.14	0.10	0.02	0.00		

Estimation on late supergames out of 10 strategies: AD, AC, Grim, TFT, STFT, T8-T6, Grim2, and TF2T.
 Rows, top 6 *played* strategies. Columns, top 7 *believed* strategies.
 Estimates in *[square brackets]* are not estimated due to collinearity.
 SFEM estimate for β is 0.94. Complete results in Table 16.

put 0 weight on T7, whereas the T7 types believe 43% of others play T7. In the Indefinite game, TFT types believe only 28% of subjects play Grim, whereas Grim types expect 80% to be Grim players.²⁴


Result 5 Beliefs are different between the Finite and Indefinite games: subjects correctly anticipate the most popular class of strategies to be different between the games (threshold vs. conditionally cooperative).

Tables 2 and 3 also reveal heterogeneity in beliefs within each game: subjects using different strategies hold different beliefs. In addition, Figure 29 in Online Appendix B shows that, in both the Finite and Indefinite games, on average subjects display a tendency to believe others are more like themselves than they actually are; that is, they overestimate the likelihood that others are of their own type.²⁵

²⁴ Estimates of $\tilde{\beta}$ close to one for all types is a results of belief reports often being extreme (close to or exactly 0 or 1), if subjects believed others implemented their strategies with error, belief reports would move towards 0.5.

²⁵This tendency can also generate large biases in beliefs for some individual types even if beliefs are fairly accurate when aggregated over types. For instance, AD types in the Indefinite game incorrectly believe others are highly likely to be AD types. These results relate to evidence from psychology and economics on the tendency to believe others act similarly to us: the false consensus effect [Ross et al., 1977]. See Blanco et al. [2014] for evidence in an experiment that elicits beliefs.

Finite



The strategy corresponding to the type is highlighted in dark grey. Analysis uses normalized stage-game payoffs.

Figure 6: Best Response for Top 6 Types in the Finite Game


Studying Best Response

Next, we explore the extent to which subjects are subjectively rational. For the purposes of our discussion, we consider *subjective rationality* in the constrained sense and examine if a subject's strategy choice is a best response (or close to being a best response) to their supgame beliefs *within* the set of strategies Z in the consideration set.²⁶ It is important to reiterate that our analysis poses no restrictions on the link between the strategies and beliefs: the strategy estimation is based on the subjects' actions and is done separately from the belief estimation, which is based on their round belief reports.

The results, presented in Figures 6 and 7, suggest most subjects' strategy choices

²⁶For consistency, the best-response analysis incorporates beliefs over implementation noise in how others carry out their intended strategy (captured by $1 - \hat{\beta}$). However, because estimated values for $\hat{\beta}$ are very close to 1, incorporating $\hat{\beta}$ does not affect the results. To calculate the expected payoff of each strategy, we simulate play in 1,000 supgames given $\hat{\beta}$.

Indefinite



The strategy corresponding to the type is highlighted in dark grey. Analysis uses normalized stage-game payoffs.

Figure 7: Best Response for Top 6 Types in the Indefinite Game

are either exact or approximate best responses given their supergame beliefs (see Online Appendix B for more). The Figures show the normalized expected payoffs (between 0, joint defection, and 1, joint cooperation) given the beliefs on the y-axis. Each bar is for one of the 10 strategies, with the one selected by that type in a darker shade of gray. In the Finite game, T7 and T6 types (38% of the population) exactly best respond to their supergame beliefs, and T8, TFT, and Grim types (39% of the population) approximately best respond to their supergame beliefs by obtaining 90%, 86%, and 89% of their best-response payoff, respectively. Of the most common six types, the only type whose strategy is far from a best response is AD (12%). In fact, their strategy choice is close to being the worst given the stated beliefs.²⁷

In the Indefinite game, a similar pattern emerges. Most common types (TFT,

²⁷Note subjects playing AD receive weakly higher payoffs in any supergame than their opponent, and these subjects have little chance to observe what would happen along alternative histories. This may contribute to why they fail to optimize given their beliefs.

Grim, Grim2, TF2T, and AD—84% of subjects) almost exactly best respond to their beliefs. One “major” type far from best responding to their belief is AC (11%), who selects the worst strategy given their beliefs. Indeed, given their beliefs, the best-response strategy is AD. For these subjects, however, some form of other-regarding preferences could reconcile strategy choices and beliefs.

Result 6 Substantial heterogeneity exists in beliefs within each game: subjects using different strategies hold different beliefs. Most types are close to best responding to their beliefs: they are subjectively rational.

Note the best-response analysis reported so far is subjective in the sense that it is based on the expected payoffs given the subjective beliefs of each type. To provide a contrast, we replicate the best-response analysis using *objective* expected payoffs computed from the strategy distribution estimated at the population level by SFEM. This analysis reveals that T6 is the best response to the population in the Finite game, and Grim2 is the best response to the population in the Indefinite game. In the Finite game, the most frequent T7 type achieves 97% of the best-response payoff from T6. In the Indefinite game, the most frequent TFT type achieves 94% of the best-response payoff from Grim2. However, some strategy-types are further away from best responding to the population. For example, the AD type in the Finite game only achieves 64% of the best-response payoff.

6 A Model of Heterogeneous Beliefs

This section presents a stylized model that generates two key results from Section 4.2 — late over-optimism in the Finite game and early over-pessimism in the Indefinite game — within a common framework. The model builds on two assumptions that are motivated by findings from the previous section: (i) Players have heterogeneous beliefs about the strategic sophistication of others, and (ii) most types best-respond to their beliefs. Specifically, the model is built on the level- k models [Stahl and Wilson, 1994, Nagel, 1995, Stahl and Wilson, 1995, Camerer et al., 2004, Costa-Gomes et al., 2001] as well as on the classic *gang of four* model [Kreps et al., 1982] as described below.²⁸

There exists a unit mass of players, and each player belongs to one of three sophistication levels $k \in \{0, 1, 2\}$. The share of level- k in the population equals ζ_k , and players are randomly matched to play the repeated PD games without observing

²⁸Gill and Rosokha [2023], in an earlier (2020) version of their paper on indefinite PDs, propose a different level- k type model to link variation in beliefs to levels of strategic sophistication.

the sophistication level of the matched player. If we denote by RD a stationary supergame strategy, which randomizes between C and D with equal probabilities in every round independent of history, the prior supergame belief of each level is specified as follows: Level-0 places probability one on RD; Level-1 places probability one on the (mixed) strategy played by level-0; Level-2 places probability $\frac{\zeta_0}{\zeta_0+\zeta_1}$ on the strategy played by level-0, and probability $\frac{\zeta_1}{\zeta_0+\zeta_1}$ on the strategy played by level-1. As for their supergame strategies, level-0 plays a mixture of Grim and RD. Level-1 and level-2 both play Grim in the Indefinite game, whereas level-1 plays T8 and level-2 plays T7 in the Finite game. Level-0 hence does not best respond to his belief, but the strategies played by level-1 and level-2 are best responses to their beliefs under permissive conditions (see Online Appendix D).²⁹ In our model, the conditionally cooperative strategy Grim played by level-0 corresponds to the irrational type who plays TFT in Kreps et al. [1982] and induces level-1 and level-2 to play cooperatively (at least initially in the Finite game).³⁰ These strategy profiles as well as the belief profiles between randomly matched pairs of players determine the evolution of the mean cooperation rates \bar{x}^t and mean round beliefs $\bar{\mu}^t$ in the population.




Figure 8: Cooperation Rates in the Finite and Indefinite Games

Notes: Average cooperation rates \bar{x}^t are shown as a solid line, and average round beliefs $\bar{\mu}^t$ are given by a dashed line. The figure is generated when $(\zeta_0, \zeta_1, \zeta_2) = (0.2, 0.5, 0.3)$ and level-0 plays Grim with probability 0.6 and RD with probability 0.4.

Figure 8 plots the evolution of the mean cooperation rates \bar{x}^t and mean round beliefs $\bar{\mu}^t$ for one (common) parameterization of the model in the Finite and Indefinite

²⁹Appendix D also illustrates how the main predictions are robust to alternative specifications of the model where (1) level-0 mixes between Grim and AD; (2) level-1 and level-2 place positive belief weight on their own level.

³⁰Although Grim (mixed with RD) by level-0 substantially simplifies the analysis in the present model, we expect the qualitative conclusion to hold if Grim is replaced by TFT as in Kreps et al. [1982].

games. In the Finite game, higher strategic sophistication corresponds to lower cooperativeness in later rounds since level-1 and level-2 types play T8 and T7, respectively. It follows that over-optimism of beliefs in later rounds is a consequence of players of all levels underestimating the sophistication of others. By contrast, in the Indefinite game, higher strategic sophistication corresponds to higher cooperativeness: level-0 is the only type defecting with positive probability in round one. It follows that underestimation of the sophistication of others generates over-pessimism of beliefs in early rounds. This is gradually corrected along the cooperative path as players update their beliefs and place higher weight on Grim played by their opponents. Hence, the model demonstrates how a player’s common but erroneous perception that others are less strategically sophisticated than them can generate the distinctive patterns of deviation in round beliefs both in the Finite and Indefinite games.

Our model builds on Kreps et al. [1982] in that it assumes the presence of the type committed to a particular conditionally cooperative strategy. The fact that the model predictions here replicate our experimental findings lends empirical support to the insight of Kreps et al. [1982] that cooperation can result from beliefs that place positive weight on such a type. The level- k structure of the present model captures the heterogeneity of beliefs among subjects as well as the discrepancy between their beliefs and strategies in a way consistent with our findings in both the Finite and Indefinite games.³¹

Online Appendix D also uses the model in this section to further highlight the empirical relevance of the presence of multiple sophistication levels. Specifically, we generate the experimental findings of Kagel and McGee [2016] and Cooper and Kagel [2023] who examine cooperation rates in the indefinitely and finitely repeated PD games when each player is replaced by a team of two players. Specifically, when we form a team by randomly matching two players and adopt the “Truth-Wins norm” by assuming that the sophistication level of a team equals that of its member with the higher sophistication level, the model replicates the evolution of cooperation rates achieved by experienced subjects in Kagel and McGee [2016] and Cooper and Kagel [2023]. Specifically, compared with individual play, the cooperation rates under team play are higher in the Indefinite game, but more accentuated in the Finite game in the sense that they are initially higher but go down more quickly and eventually become lower in later rounds of the Finite game.

³¹It is possible to interpret the reputation model of Kreps et al. [1982] as a level- k model in which (1) there exist two levels of sophistication: the irrational level-0 and the rational level-1, and (2) level-1’s belief is correct and places large weight on the level-1 strategy. A critical difference then is that the players’ beliefs in the present model are misspecified.

7 Robustness

Section 7.1 provides analyses on the robustness of the estimation method we adopt to recover beliefs over strategies. Section 7.2 presents results from two new treatments that study beliefs in the Indefinite game with different stage-game payoffs.

7.1 Estimation of Supergame Strategies

In Online Appendix E.1, we report results of simulations. Those show that supergame beliefs can be accurately recovered in data sets similar to ours by using the belief estimation procedure described in this paper. Below, we explain how the procedure can be generalized to be used with other methods to type subjects, allow for non-Bayesian updating, and how such changes impact our main results.

Typing in stage (A), as performed in this paper, is potentially impacted by two factors: (i) the consistency of actions with each strategy and (ii) the prior likelihood of each strategy. One potential concern is that the prior likelihood of each strategy can have a disproportionate impact, distorting typing such that subjects end up being classified as using strategies that are popular at the population-level even though their actions are not closely consistent with the strategy. Such a concern is not warranted in our data set: The typing procedure assigns *all but one* subject to the strategy their actions are most consistent with according to (i).³² Namely, in our data set, incorporating (ii), the prior likelihood of each strategy into the typing procedure impacts results mostly as a tie-breaking rule, allowing us to uniquely classify subjects whose actions are equally consistent with multiple strategies.³³

A related second potential concern is that by using a typing method that generates unique classification, we are also estimating beliefs for subjects who are not well identified in terms of their strategy choice (e.g., those whose actions are equally consistent with multiple strategies). To respond to this concern, we study the extent to which our belief estimates change if we restrict our analysis to only those subjects who are well differentiated in terms of their strategy choice. Results from such an exercise are presented in Online Appendix E.2. In summary, while different typing

³² The only exception is a subject in the Indefinite game who took 17 actions (out of 24) consistent with T6, but 16 actions consistent with Grim. Because the SFEM estimate for T6 is less than 1 percent in this treatment, the Bayesian posterior that the subject is playing Grim is higher than that of T6, and thus the subject is typed as Grim rather than T6.

³³ Simulations (reported in Online Appendix E.1) demonstrate that beliefs can be reliably recovered when using this typing procedure as part of the estimation approach. Online Appendix E.2 provides further analysis on the ability of this procedure to differentiate between types.

methods produce slightly different beliefs estimates, the main patterns echo those in our main analysis.

A third potential concern with the method we propose is that it assumes Bayesian updating. This simplifies the conceptual framework and serves as a reasonable benchmark. Nonetheless, the method can be generalized to incorporate non-Bayesian updating. In Online Appendix E.3 we conduct such an analysis to confirm that the main results of the paper are robust to allowing for such behavior.

7.2 Different Parameters Within the Same Game Structure

Our results so far establish that beliefs (both about actions and supergame strategies) capture the key differences in strategic behavior between the Finite and Indefinite games. That is, keeping the stage game constant, subjects' beliefs change as we vary the termination rule. This section investigates a complementary question: keeping the game structure (termination rule) constant, how do beliefs change with the stage game parameters?

To study this, we focus on the Indefinite game, and conduct two additional treatments that preserve the same $\delta = 7/8$, but vary how conducive stage-game parameters are to cooperation. Below we present results from 16 new sessions: 8 where the temptation payoff of the stage game is increased to 73 (from an original value of 63, referred to as the *High T* treatment) and 8 where the reward to joint cooperation is decreased to 45 (from original value of 51, referred to as the *Low R* treatment). Prior literature suggests cooperation to be more challenging in these new treatments, but more so in the *Low R* than *High T* treatment.³⁴ Further details on the implementation of these treatments are provided in Online Appendix E.4.

Figure 9 shows average beliefs and average probability of cooperation in the original Indefinite game and contrasts these to the new treatments. As can be seen, cooperation rates in the *High T* treatment are quite similar to those from the original game. Beliefs are slightly more pessimistic about the likelihood of cooperation in round one ($p = 0.064$). However, cooperation rates mark a large decrease in the *Low R* treatment relative to the original game. Consistent with this, subjects expect cooperation to be less likely ($p < 0.001$ for both actions and beliefs in round one and the first eight rounds altogether).

In both new treatments, as in the original Indefinite game, beliefs in round two,

³⁴Previous experiments with parameters comparable to *High T* display large variations in cooperation rates (see Dal Bó and Fréchette [2018]).





Figure 9: Choices and Beliefs by Round (Equivalent to Figure 2)

conditioning on history, show large movements from round one beliefs; these movements are in the correct direction in both the *High T* and *Low R* treatments (see Figures 49 and 50 in Online Appendix E.4). In the *Low R* treatment, we not only see large adjustments downwards (as in the original game), but also substantial adjustment upward in the case where both subjects cooperate in round one (different from the original game). Online Appendix E.4 establishes that other key findings from the original game are also replicated in these new treatments.

The distribution of strategies and the beliefs over strategies further reflect how changing the stage game parameters impacts strategic reasoning in the Indefinite game. Figure 10 depicts the cumulative distribution of strategies and supergame beliefs in the new treatments (*High T* and *Low R*) and contrast these with those from the original Indefinite game (see Online Appendix E.4). Strategies are ranked in terms of their *cooperativeness*. Formally, we define a strategy to be more cooperative than another one if, as the probability of implementation errors goes to zero (i.e. as $\beta \rightarrow 1$), the expected payoff associated with playing the former strategy against itself is higher than the expected payoff of playing the latter strategy against itself (derivation provided in Online Appendix G).³⁵ Given the cooperativeness ranking, first order stochastic dominance between distributions can be interpreted as a treat-

³⁵On the subset of strategies considered by Proto et al. [2020], our cooperativeness order coincides with the inverse of their *harshness* ranking.




Late supergames.

Figure 10: Distribution of Strategies and Average Beliefs over Strategies in the Indefinite Treatments

ment shifting behavior to be less cooperative (in the left panel) and beliefs to be more pessimistic about others' cooperativeness (in the right panel). Thus, Figure 10 reveals: (i) As behavior becomes less cooperative, beliefs also become more pessimistic about others' cooperativeness (as seen from the comparison of the original game to the one with *Low R*).³⁶ This could explain why beliefs are more pessimistic in that treatment in later supergames. (ii) Beliefs in the Indefinite game (particularly with parameter values that are not very conducive to cooperation) underestimate the cooperativeness of others (as seen most clearly from the comparison of the distributions for the *Low R* treatment on the left and right panels).

Finally, Figure 11 focuses on the most common four types in each of the new treatments and indicates whether their strategy choice is subjectively rational given their supergame beliefs. The darker bars depicting expected payoff associated with the strategy corresponding to the type are either maximal or close to being so in all eight panels, reaffirming our earlier results that subjects are close to best responding given their supergame beliefs. Focusing on heterogeneity within each treatment,

³⁶Beliefs are also more pessimistic in the the *High T* treatment although cooperation rates are similar to the original game. This could be driven by the differences in cooperation rates between these two treatment in early supergames as can be seen from Figure 48 in Online Appendix E.4.



Percentages refer to SFEM shares. The strategy corresponding to the type is highlighted in dark grey. Analysis uses normalized stage-game payoffs.

Figure 11: Best Response for Top 4 Types in New Indefinite Treatments

we observe the subjects typed as playing more cooperative strategies to have more optimistic beliefs about the cooperativeness of the others. This is visible in the Figure when comparing the height of the bars across types, but is more directly observable using the estimates in Tables 29 and 30. Focussing on the TF2T and AD types of the *High T* treatment, the belief estimates indicate that TF2T types believe 80% of subjects play cooperative strategies, whereas AD types only believe 56% of subjects do so. In the *Low R* treatment, Grim2 players believe 50% of subjects play cooperative strategies, versus 19% among AD types.

Overall, these results strengthen our earlier findings on heterogeneity in strategy choice and its close connection to beliefs in the Indefinite game. Despite experience with the environment, subjects hold very different beliefs about the strategy choice of their opponent in the Indefinite game. Differences in beliefs, to a large extent, support differences in strategy choice. The new treatments demonstrate this very clearly. Subjects who play cooperative and uncooperative strategies have sufficiently different beliefs such that strategy choice is subjectively rational (or close to being so) in each of the cases.

8 Conclusion

Beliefs play a central role in equilibrium theory, and increasing evidence suggests they are also key to understanding behavior observed in repeated settings. This study elicits beliefs in finitely and indefinitely repeated PD games with the main goal of providing a novel data set to inform our views on how beliefs, actions, and strategy choices are linked in this important class of games.

We separate the discussion of our findings into those from round beliefs (beliefs over actions) and supergame beliefs (beliefs over strategies). Our first key finding is that round beliefs are, in aggregate, remarkably accurate. In both the Finite and Indefinite games, round beliefs averaged over all rounds are less than three percentage points away from the empirical action frequencies in our main treatments. Round beliefs also adjust appropriately to the history of play even when these adjustments are not small: in some histories, they move by more than 50 percentage points between rounds one and two. However, there are small, but systematic deviations: over-optimism in late rounds of the Finite game and over-pessimism in early rounds of the Indefinite game. In addition, the early over-pessimism observed in the Indefinite game is also confirmed in two additional treatments. Another key finding is that beliefs over stage actions are *forward looking*. Most notably, beliefs along the history of mutual cooperation evolve very differently in the Finite and Indefinite games. Persistence of cooperation in the Indefinite game and its collapse late in the Finite game are correctly anticipated along such histories. Interestingly, the same action choice can be observed in both games even when subjects report very different beliefs.

Our second category of findings is based on the development of a novel method to recover supergame beliefs from the sequence of round beliefs. These supergame beliefs correctly capture the different classes of strategies used in each environment—threshold strategies in the Finite game and conditionally cooperative strategies in the Indefinite game—and also display substantial heterogeneity across subjects playing different strategies. This heterogeneity in strategies can be linked to the heterogeneity in supergame beliefs as most types are close to being subjectively rational: given their beliefs, their selected strategy is optimal (or close to it) among the strategies considered. Although beliefs are surprisingly accurate as noted above, systematic deviations at key junctures of the game can have important implications for behavior. In the Finite game, subjects tend to believe others play more cooperative strategies than their own, which can slow down unravelling. This is consistent with the findings from Kagel and McGee [2016] where team-dialogues reveal subjects engage in limited backward induction and fail to account for others reasoning in a similar way. In the Indefinite game, as particularly evident in our additional treatments, subjects

believe others use defective strategies more than is the case. This can explain why payoffs observed in experiments on the indefinitely repeated PD games are often far below the maximum symmetric equilibrium payoffs.

The procedure proposed here to recover supergame beliefs has broad applicability. In repeated games, it can be applied to different sets of strategies and/or be combined with alternative methods to type subjects. More generally, this procedure can be used to recover beliefs over strategies from beliefs over actions in any sequential game. Our aim in eliciting round beliefs is to keep belief elicitation in the laboratory simple and non-invasive so as to minimize the impact of belief elicitation on behavior. Although the method we use to recover supergame beliefs inevitably makes assumptions on how round and supergame beliefs are linked, this approach is a useful starting point and, as discussed further in the paper, can be modified when there are concerns about the suitability of those assumptions.

Our results also provide insights into the forces that underlie some of the key behavioral patterns observed in these games. In particular, they show that standard preferences with optimizing behavior under slightly erroneous beliefs go a long way to account for the observed behavior in both the Finite and Indefinite games. In the finitely repeated PD games, small but systematic departures from accurate beliefs (at key points in the supergame) sustain cooperation. Although beliefs are generally accurate, for 80% of subjects in the Finite game, best responding to their *subjective* beliefs (that are slightly over-optimistic) involves cooperating more than would be *objectively* optimal against the actual strategy distribution in the population. In the indefinitely repeated PD games, our results highlight the difficulty of resolving equilibrium selection due to the persistence of heterogeneous beliefs. This is particularly true in environments that are conducive to cooperation since subjects experience few histories which prove those beliefs to be incorrect, leaving a variety of conditionally cooperative strategies popular even after many repetitions. The systematic deviations in both the Finite and Indefinite games can be replicated by a stylized model with *level-k* agents. This model brings to light the intuition that late over-optimism in the Finite game and early over-pessimism in the Indefinite game, which at first glance appear at odds, can result from a common mistake where players believe others to be less strategically sophisticated than themselves.

In summary, our results on beliefs suggest subjects understand the different consequences of the finite and indefinite horizons even when their observed behavior is identical in early rounds of the repeated games. In other words, subjects have a refined awareness of the rules of the game and the implications of these rules for the dynamics of cooperative behavior. They also suggest the calculus underpinning

choices are very different across finitely and indefinitely repeated environments.

References

- Franklin Allen. Discovering personal probabilities when utility functions are unknown. *Management Science*, 33(4):542—544, 1987.
- Ned Augenblick, Eben Lazarus, and Michael Thaler. Overinference from weak signals and underinference from strong signals. *arXiv preprint arXiv:2109.09871*, 2023.
- Cuimin Ba, J Aislinn Bohren, and Alex Imas. Over-and underreaction to information. *Available at SSRN*, 2023.
- Dan Benjamin, Aaron Bodoh-Creed, and Matthew Rabin. Base-rate neglect: Foundations and implications. *Base-Rate Neglect: Foundations and Implications*, 62, 2019.
- Daniel J Benjamin. Errors in probabilistic reasoning and judgment biases. *Handbook of Behavioral Economics: Applications and Foundations 1*, 2:69–186, 2019.
- Mariana Blanco, Dirk Engelmann, Alexander K Koch, and Hans-Theo Normann. Preferences and beliefs in a sequential social dilemma: a within-subjects analysis. *Games and Economic Behavior*, 87:122–135, 2014.
- Colin F Camerer, Teck-Hua Ho, and Juin-Kuan Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3):861–898, 2004.
- Yin-Wong Cheung and Daniel Friedman. Individual learning in normal form games: Some laboratory results. *Games and Economic Behavior*, 19(1):46 – 76, 1997.
- David J Cooper and John H Kagel. Using team discussions to understand behavior in indefinitely repeated prisoner’s dilemma games. *American Economic Journal: Microeconomics*, 15:114–145, 2023.
- Miguel Costa-Gomes, Vincent P Crawford, and Bruno Broseta. Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5):1193–1235, 2001.
- Miguel A. Costa-Gomes and Georg Weizsäcker. Stated Beliefs and Play in Normal-Form Games. *The Review of Economic Studies*, 75(3):729–762, 07 2008.
- Paolo Crosetto and Antonio Filippin. The “bomb” risk elicitation task. *Journal of Risk and Uncertainty*, 47(1):31–65, 2013.

- Pedro Dal Bó. Cooperation under the shadow of the future: experimental evidence from infinitely repeated games. *American Economics Review*, 95(5):1591–1604, December 2005.
- Pedro Dal Bó and G.R. Fréchette. The evolution of cooperation in infinitely repeated games: Experimental evidence. *American Economic Review*, 101(1):411–429, 2011.
- Pedro Dal Bó and Guillaume R Fréchette. On the determinants of cooperation in infinitely repeated games: A survey. *Journal of Economic Literature*, 56(1):60–114, 2018.
- Pedro Dal Bó and Guillaume R. Fréchette. Strategy choice in the infinitely repeated prisoners dilemma. *American Economic Review*, 2019. Forthcoming.
- David Danz, Lise Vesterlund, and Alistair J Wilson. Belief elicitation: Limiting truth telling with information on incentives. Technical report, National Bureau of Economic Research, 2020.
- David N. Danz, Dietmar Fehr, and Dorothea Kübler. Information and beliefs in a repeated normal-form game. *Experimental Economics*, 15(4):622–640, Dec 2012.
- Douglas Davis, Asen Ivanov, and Oleg Korenok. Individual characteristics and behavior in repeated games: An experimental study. *Experimental Economics*, 19(1):67–99, 2016.
- John Duffy, Ed Hopkins, and Tatiana Kornienko. Facing the grim truth: Repeated prisoner’s dilemma against robot opponents. *Working Paper*, 2021.
- Fabian Dvorak. stratEst: Strategy estimation in R. *Dvorak, F.(2020). stratEst: Strategy Estimation in R. TWI Working Paper Series*, (119), 2020.
- Matthew Embrey, Guillaume R Fréchette, and Sevgi Yuksel. Cooperation in the finitely repeated prisoner’s dilemma. *The Quarterly Journal of Economics*, 133(1):509–551, 2018.
- Ignacio Esponda, Emanuel Vespa, and Sevgi Yuksel. Mental models and learning: The case of base-rate neglect. *Working paper*, 2022.
- Urs Fischbacher. z-Tree: Zurich toolbox for readymade economic experiments. *Experimental Economics*, 10(2):171–178, 2007.

- Urs Fischbacher and Simon Gächter. Social preferences, beliefs, and the dynamics of free riding in public goods experiments. *American Economic Review*, 100(1): 541–56, March 2010.
- Guillaume R. Fréchette. Session-effects in the laboratory. *Experimental Economics*, 15(3):485–498, September 2012.
- Guillaume R. Fréchette and Sevgi Yuksel. Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination. *Experimental Economics*, 20(2):279–308, 2017.
- Drew Fudenberg, David G. Rand, and Anna Dreber. Slow to anger and fast to forget: Leniency and forgiveness in an uncertain world. *American Economic Review*, 102(2):720–749, 2012.
- Simon Gächter and Elke Renner. The effects of (incentivized) belief elicitation in public goods experiments. *Experimental Economics*, 13(3):364–377, September 2010.
- David Gill and Yaroslav Rosokha. Beliefs, learning, and personality in the indefinitely repeated prisoner’s dilemma. Working paper, 2020.
- David Gill and Yaroslav Rosokha. Beliefs, learning, and personality in the indefinitely repeated prisoner’s dilemma. *American Economic Journal: Microeconomics* (Forthcoming), 2023.
- Ben Greiner. Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125, 2015.
- David M Grether. Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly journal of economics*, 95(3):537–557, 1980.
- Tanjim Hossain and Ryo Okui. The binarized scoring rule. *Review of Economic Studies*, 80(3):984–1001, 2013.
- Kyle Hyndman, Antoine Terracol, and Jonathan Vaksman. Beliefs and (in) stability in normal-form games. *Experimental Economics*, 25(4):1146–1172, 2022.
- John H Kagel and Peter McGee. Team versus individual play in finitely repeated prisoner dilemma games. *American economic Journal: Microeconomics*, 8(2):253–76, 2016.

- Daniel Kahneman and Amos Tversky. On the psychology of prediction. *Psychological review*, 80(4):237, 1973.
- Ehud Kalai and Ehud Lehrer. Rational learning leads to nash equilibrium. *Econometrica*, 61(5):1019–1045, 1993. ISSN 00129682, 14680262.
- David M. Kreps, Paul Milgrom, John Roberts, and Robert Wilson. Rational cooperation in the finitely repeated prisoners’ dilemma. *Journal of Economic Theory*, 27(2):245–252, 1982.
- Volodymyr Lugovskyy, Daniela Puzzello, and James Walker. On cooperation in finitely and indefinitely repeated prisoner’s dilemma games. *Working Paper*, 2020.
- Richard D McKelvey and Talbot Page. Public and private information: An experimental study of information pooling. *Econometrica: Journal of the Econometric Society*, pages 1321–1339, 1990.
- John H Nachbar. Beliefs in repeated games. *Econometrica*, 73(2):459–480, 2005.
- Rosemarie Nagel. Unraveling in guessing games: An experimental study. *The American economic review*, 85(5):1313–1326, 1995.
- Tibor Neugebauer, Javier Perote, Ulrich Schmidt, and Malte Loos. Selfish-biased conditional cooperation: On the decline of contributions in repeated public goods experiments. *Journal of Economic Psychology*, 30(1):52 – 60, 2009.
- Yaw Nyarko and Andrew Schotter. An experimental study of belief learning using elicited beliefs. *Econometrica*, 70(3):971–1005, 2002.
- Eugenio Proto, Aldo Rustichini, and Andis Sofianos. Intelligence personality and gains from cooperation in repeated interactions. *Journal of Political Economy*, 127(3):1351–1390, 2019.
- Eugenio Proto, Aldo Rustichini, and Andis Sofianos. Intelligence, errors and strategic choices in the repeated prisoners’ dilemma. *Working Paper*, 2020.
- Eugenio Proto, Aldo Rustichini, and Andis Sofianos. Intelligence, errors, and cooperation in repeated interactions. *The Review of Economic Studies*, 89(5):2723–2767, 2022.
- Pedro Rey-Biel. Equilibrium play and best response to (stated) beliefs in normal form games. *Games and Economic Behavior*, 65(2):572 – 585, 2009. ISSN 0899-8256.

- Julian Romero and Yaroslav Rosokha. Constructing strategies in the indefinitely repeated prisoner’s dilemma game. *European Economic Review*, 104:185–219, 2018.
- Julian Romero and Yaroslav Rosokha. Mixed strategies in the indefinitely repeated prisoner’s dilemma. Working Paper, 2023.
- Lee Ross, David Greene, and Pamela House. The “false consensus effect”: An egocentric bias in social perception and attribution processes. *Journal of experimental social psychology*, 13(3):279–301, 1977.
- Alvin E. Roth and J. Keith Murnighan. Equilibrium behavior and repeated play of the prisoner’s dilemma. *Journal of Mathematical Psychology*, 17:189–198, 1978.
- Karl H. Schlag and Joël J. van der Weele. Eliciting probabilities, means, medians, variances and covariances without assuming risk neutrality. *Theoretical Economics Letters*, (3):38–42, 2013.
- Dale O Stahl and Paul W Wilson. Experimental evidence on players’ models of other players. *Journal of economic behavior & organization*, 25(3):309–327, 1994.
- Dale O Stahl and Paul W Wilson. On players models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254, 1995.
- Emanuel Vespa and Alistair Wilson. Experimenting with the transition rule in dynamic games. *Quantitative Economics*, 10:1825–1849, 2019.
- Alistair J. Wilson and Emanuel Vespa. Paired-uniform scoring: Implementing a binarized scoring rule with non-mathematical language. Working Paper, 2018.