ISSN (Print) 0473-453X
Discussion Paper No. 1119
ISSN (Online) 2435-0982

BELIEFS IN REPEATED GAMES

Masaki Aoyagi
Guillaume R. Fréchette
Sevgi Yuksel

Secondly Revised May 2022
Revised November 2021
February 2021

The Institute of Social and Economic Research
Osaka University
6-1 Mihogaoka, Ibaraki, Osaka 567-0047, Japan

Beliefs in Repeated Games*

Masaki Aoyagi
Osaka University

Guillaume R. Fréchette
NYU

Sevgi Yuksel
UCSB

May 9, 2022

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' beliefs about the other player's action are generally accurate despite some small systematic deviations corresponding to early pessimism in the indefinitely repeated game and late optimism in the finitely repeated game. The data reveals a close link between beliefs and actions that differs between the two games. In particular, the same history of play is associated with different beliefs, and the same beliefs with different action choices in each game. Moreover, we find beliefs anticipate the evolution of behavior within a supergame, changing in response to the history of play (in both games) and the number of rounds played (in the finitely repeated game). We then use the subjects' beliefs over actions in each round to identify their beliefs over supergame strategies played by the other player. We find these beliefs correctly capture the different classes of strategies used in each game. Importantly, subjects using different strategies have different beliefs, and for the most part, strategies are subjectively rational given beliefs. The results also suggest subjects tend to underestimate the likelihood that others use less cooperative strategies. In the finitely repeated game, this helps explain the slow unravelling of cooperation. In the indefinitely repeated game, persistence of heterogeneity in beliefs underpins the difficulty of resolving equilibrium selection.

JEL classification: C72, C73, C92.
Keywords: repeated game, belief, strategy, elicitation, prisoner's dilemma.

*We wish to thank Hazem Alshaikhmubarak for excellent research assistance. We would further like to thank the 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; Fréchette: frechette@nyu.edu; Yuksel: sevgi.yuksel@ucsb.edu.

1 Introduction

Social dilemmas encompass a large class of situations of much interest in the social sciences. Examples in economics are numerous, ranging from Cournot competition to natural resource extraction. Among them, the prisoner's dilemma (PD) captures in its simplest form a tension between individual payoff maximization and social efficiency. How this tension is resolved in a repeated setting as a function of environmental parameters—payoffs, monitoring technology, discounting, the game horizon, and so on—has been an active area of theoretical research. However, our empirical understanding of behavior in repeated games is much more limited. Although the number of experiments on repeated games is increasing, the bulk of our knowledge concerns how the level of cooperation varies with the environmental parameters; and many facets of the documented behavior are still black boxes. For example, we still have a limited understanding of the degree to which frequently observed patterns of behavior (such as sustained cooperation in a finitely repeated PD) are driven by non-standard preferences, incorrect beliefs, or bounded rationality. In this paper, by making beliefs observable, we bring to light one key force to help us better understand behavior in such games.

A player's beliefs about other players' strategies form the foundation of equilibrium analysis: beliefs are assumed to correctly identify the strategies played by other players, and the strategies best respond to those beliefs. It is surprising that the validity of these assumptions have not been tested in a game as canonical as the repeated PD. To the extent that beliefs have been studied experimentally in games, the focus has been on the consistency of beliefs and actions in 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], focussing 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, concludes that there may be no general findings in this regard, as it may depend on features of the game, but their results suggest that best-response behavior may be higher in simple games. There are reasons to suspect that these results might not extend to repeated games. First, in repeated games, strategies can be very complex because they are complete contingent plans that specify actions after every history, and as in the case of repeated PD, many strategies can be rationalized as best responses to some beliefs. Second, multiplicity of equilibria in repeated games further complicates forming correct beliefs about others. Third, while our focus is on experienced behavior after subjects have played multiple supergames, learning is made difficult given the extremely large number of histories that the subjects may

experience. We provide feedback to the subjects, which has been shown to increase the fraction of best response behavior in one-shot games (Danz et al. [2012], for instance). However, at least one third of the subjects experience a history that is new to them in their final supergame. As such, a subject cannot simply use past experiences to form beliefs.[1]

The theoretical contrast between the finitely and indefinitely repeated PD provides a useful backdrop for the study of beliefs and their relationship to cooperation. The unique equilibrium entails no cooperation in the finitely repeated PD, but a multitude of outcomes ranging from no cooperation to full cooperation are compatible with equilibrium behavior in the infinitely repeated PD for sufficiently patient players. As such, our experiment consists of two treatments: the *Finite* game and the *Indefinite* game. In the former, subjects play eight rounds of a PD; in the latter, subjects play PDs over a random number of rounds with a continuation probability of 7/8.[2] We selected these parameters and the stage game based on prior results in the literature: they are expected to generate not only significant levels of cooperation in both the Finite and Indefinite games, but also similar levels of round-one cooperation in both environments. Finding parameters that would generate very different initial cooperation rates between these two games is easy [Dal Bó, 2005]. In contrast, we intended to create two treatments where behavior was expected to be similar despite the theoretical difference. The treatment variation hence permits a comparison of beliefs of subjects taking the same action in the same round (potentially along the same history) across the Finite and Indefinite games, and provides insight into whether their strategic reasoning is similar or different across these two games.

In a first foray into beliefs in repeated PD games, many questions could be 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. For example, we use games with perfect monitoring where the past actions of both players are observed without noise, instead of games with imperfect (public or private) monitoring. More importantly, we only elicit (first-order) beliefs about the other players' stage actions and not, for example, beliefs conditional on some

[1] In fact, even in simpler games, Nyarko and Schotter [2002] find that beliefs are not simply a summary of past experiences, or more precisely they are not equivalent to γ-weighted empirical beliefs [Cheung and Friedman, 1997]. That is, consistent with our results, as we find that subjects' beliefs are not equivalent to the frequency of past actions.

[2] Indefinite repetition (first introduced by Murnighan and Roth [1983]) is the standard method of implementing infinitely repeated games in the laboratory. The continuation probability of the indefinite game is associated with the discount factor of the infinitely repeated game.

action realization, beliefs over beliefs, or beliefs over supergame strategies. Despite their simplicity round beliefs already reveal important information. For one thing, their accuracy is a necessary condition for subjects to correctly predict other players' strategies. They also inform us about what subjects anticipate about other's action and how this is informed by past experiences, history of play, as well as other features of the game. These observations provide important insights on how subjects strategically reason in repeated games. For instance, we can investigate whether beliefs evolve differently—along identical histories—in a finitely vs. indefinitely repeated PD, allowing us to see directly whether subjects strategically differentiate between these two games even when their actions do not reflect these differences.

To go beyond round beliefs, we propose a novel method to recover supergame beliefs. The method is more general than the application presented here and can be used more broadly to construct ex-ante beliefs over extensive form games when the beliefs are only elicited at realized nodes. This can be useful for two reasons: (1) when eliciting beliefs over all potential histories is not possible; or (2) when there are reasonable concerns that an experimental design that allows for such belief elicitation might alter how subjects approach the strategic interaction. In our context, the approach we adopt involves two steps: First, we type subjects according to the supergame strategy they are estimated to be playing among a predefined set. Second, we estimate, for each subject type separately, the supergame beliefs over this set.[3]

The previous experimental literature on repeated games provides little guidance on what to expect with respect to consistency of beliefs and actions in this environment. Given the complexity of repeated games, and the large number of potential histories, whether or not subjects will have accurate beliefs (even after gaining experience) and whether subjects will best respond to their beliefs is unclear. It is also not obvious whether one would expect beliefs to be more accurate and behavior to be more consistent with best response in the Finite game, in which there is a unique equilibrium, but observed behavior is not stationary, or the Indefinite game, which is marked by equilibrium multiplicity. Furthermore, should we expect beliefs to be more accurate and consistent with actions in early rounds of a supergame where the subjects find most histories familiar or in later rounds where they become more informed about their opponent?

We identify three classes of key results. First, beliefs are, broadly speaking, *ac-*

[3]This formulation is close to that explored in Kalai and Lehrer [1993], who show that if players of an infinitely repeated game start with subjective beliefs about the opponents' strategies that place positive probability on their true strategies, Bayesian updating will lead in the long run to the NE play of the repeated game.

curate. This is noticeable at many levels: Round by round, unconditional average beliefs are close to empirical action frequencies. 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. However, history of play doesn't capture all the variation in beliefs: Subjects experience at least 156 histories in a supergame in each treatment. As a result, many specific histories are completely new (as mentioned earlier, more than a third of subjects experience a history that is new to them in the final supergame). As such, subjects cannot simply rely on past experiences to form beliefs. 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. Beliefs are informed by past experiences and subjects' own strategic considerations, but cannot be reduced to either. For instance, in more than three of quarters of cases, subjects' round one beliefs differ from the cooperation rate they have experienced 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 experienced cooperation rates.

Second, despite the aforementioned general accuracy, beliefs also display small but systematic deviations. Beliefs are too optimistic at the end of the Finite game and too pessimistic at the beginning of the Indefinite game. Overall, beliefs correctly anticipate the types of strategies played in each environment: conditionally cooperative strategies that are stationary in the Indefinite game, but switch to defection in the last few rounds in the Finite game. Nonetheless, such beliefs are not necessarily perfectly calibrated to the actual frequency of strategies in the population. In the case of the Finite game this plays a key role in slowing down unravelling of cooperation. In the Indefinite game, this may be part of the explanation for the fact that payoffs are inside the efficient frontier.

Third, belief heterogeneity is important. This heterogeneity is directly visible in round one beliefs: thus indicating that subjects expectations about whether others are more or less likely to use cooperative versus defective strategies vary. It is also present in the estimated beliefs over strategies that display important variation across types. Moreover, for most types, this variation in beliefs helps rationalize subjects'

decisions to cooperate or not: given beliefs, their strategy choice is optimal (or close to optimal) among the strategies considered.

How do these findings inform our understanding of behavior in these games? In our finitely repeated PD, we observe high cooperation and slow unravelling, behavior not predicted by theory. Much of this can be explained by beliefs that are just slightly overly 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 similar to many that have been studied before), even a little overoptimism can substantially weaken incentives to defect earlier, hindering unravelling of cooperation. Our estimates suggests that, for 80% of subjects, best responding to their beliefs translates into *cooperating too much* relative to the empirical best response strategy. In the indefinitely repeated PD, instead of coordination (convergence to an equilibrium), a variety of strategies are observed to persist in the long run. We see that beliefs remain heterogeneous with experience, which is difficult to correct because the histories that distinguish the more common strategies are rarely observed. The heterogenous strategies are subjectively optimal given the heterogenous beliefs.

There is a rapidly growing literature on experiments with belief elicitation, but most of the papers in the literature examine beliefs in individual decision making settings (see Danz et al. [2020] for a recent review). Those that study beliefs in games mostly use one-shot games.[4] Gächter and Renner [2010] study how contributions in a public goods game are related to preferences and beliefs. Although they do discuss *unravelling*, a key feature of Finite games, their use of the term is different from ours (and the repeated game literature). Their focus is on the decrease in the probability of cooperation in a series of one-shot games. We are interested in the point at which players start defecting in a repeated game. To the extent that prior experiments have induced repeated games in the laboratory, they do so assuming that incentives in static interactions remain unchanged in repeated play and do not analyze the dynamic incentives of the players.[5] In contrast, our experiment elicits beliefs in repeated games where dynamic incentives are clearly important.

As mentioned earlier, evidence with respect to the accuracy of beliefs and best response behavior in one-shot games is varied (Nyarko and Schotter [2002], Costa-

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

[5]For example, some studies use games for which the equilibrium payoff set does not expand with repetition. In others, the repeated game is simply a byproduct of a design that uses fixed-pairing among subjects. As such, analyses in these papers do not address dynamic strategies, and no scope exists for subjects to learn over multiple supergames.

Gomes and Weizsäcker [2008], Rey-Biel [2009], Danz et al. [2012], Hyndman et al. [2010], Hyndman et al. [2012]). Duffy et al. [2021] find failure of best response in a design where subjects play an indefinitely repeated PD against robot players which are known to follow the Grim trigger strategy. We find that despite the complexity and richness of repeated games, subjects have fairly accurate beliefs and best-response behavior is high. This finding is in line with a closely related paper by Gill and Rosokha [2020]. However, both their focus and design are substantially different. Gill and Rosokha [2020] study indefinitely repeated (but not finitely repeated) PDs. Their subjects directly choose strategies (from a list of 10) for which they measure beliefs over those 10 strategies twice: in the first and last supergames. Hence, they focus more directly on the change in beliefs. They also link personality traits to strategies and beliefs. See Online Appendix A for a more detailed review of experimental studies with belief elicitation.

The paper is organized as follows. The formal description of strategies and beliefs are given in section 2. section 3 describes the experimental design. Results are presented in section 4. We conclude with a discussion in section 5.

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 $t = 1, 2, \ldots$, history h^t of length t is a sequence of action profiles in rounds $1, \ldots, t$. Let $H^t = A^t$ be the set of t-length histories. A player's (behavioral) *strategy* $\sigma_i = (\sigma_i^1, \sigma_i^2, \ldots)$ 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:

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

where E_σ is the expectation with respect to the probability distribution of $h^T = (a^1, \ldots, a^T)$ 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:

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

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. Specifically, we suppose player i believes j's strategy is randomly chosen from some *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*.[6] One interpretation of $\tilde{p}_i$ is that it represents i's prior belief over the proportion of different strategies played by other subjects in that session.[7]

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

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

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 4.3, we recover the subjects' supergame beliefs $\tilde{p}_i$ from the sequence of their elicited beliefs $\mu_i^1, \mu_i^2, \ldots$.

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 $h^t = (h^{t-1}, a^t)$,

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

[6] We 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.

[7] With 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.

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 $\tilde{p}_i$:

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

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 $h^{t-1} = ((C,C), \ldots, (C,C))$ and $\sigma_i^t(h^{t-1})(C) = 0$ otherwise. σ_i is TFT (resp. STFT) if $\sigma_i^1(C) = 1$ (resp. $\sigma_i^1(D) = 1$) and $\sigma_i^t(h^{t-1})(a_j^{t-1}) = 1$ for every h^{t-1} and $t \geq 2$. For $k = 1, 2, \ldots$, 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.[8]

3 Design

The experiment involves two (between-subjects) treatments, which we refer to as the *Finite* and *Indefinite* games. Three important considerations (besides the aforementioned aim for simplicity) guided 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. 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. Many other comparisons (such as changing parameters within game type) would be interesting, but they seemed to offer less potential to deliver novel insights.

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.

[8] All strategies considered in our analysis are listed and defined in Table 15 of Online Appendix B.

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 [2019], who do this for strategy elicitation.[9] 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. Indeed, our results: 1) Reproduce the qualitative features of previous experiments without belief elicitation (with similar parameters), those are documented later in Result 4. 2) 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 10). 3) In the case of the Finite game treatment, one prior experiment used exactly the same parameters (but did not elicit beliefs). Figure 13 shows that cooperation rates in the current experiment when beliefs are elicited are not different from those in that experiment where they were not.[10]

3. *Allowing subjects to gain ample experience.* Prior research, both with finitely and indefinitely PD games, show the importance of experience [Embrey et al., 2018, Dal Bó and Fréchette, 2018], in that behavior evolves in important ways and subjects need time to understand dynamic incentives.[11] 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. Asking more complex belief questions would necessarily slow down the experiment and reduce the number of supergames.

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.[12] We

[9]Dal Bó and Fréchette [2019] find that choices in their experiments with strategy elicitation (introduced after a period of play of the standard repeated PD) are similar to those from experiments without strategy elicitation. Romero and Rosokha [2019] also find that choices are unaffected with such a design. Other experiments that immediately introduce strategy elicitation have reported different results.

[10]Cooperation rates are different in supergames one and two for Round one cooperation and in supergame one for Cooperation in All Rounds. We attribute those differences to small sample variations or subject pool effects because beliefs are not elicited in either case. However, those differences disappear quickly.

[11]Whether experience should be defined in terms of the number of supergames or of the number of total rounds across multiple supergames is not clear.

[12]The normalization facilitates comparison with prior studies. With normalization, we set the

Table 1: Stage Game

In ECU

	C	D
C	51, 51	22, 63
D	63, 22	39, 39

Normalized

	C	D
C	1, 1	$-1.41\overline{6}$, 2
D	2, $-1.41\overline{6}$	0, 0

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, $T = 8$.[13] In the Indefinite game, after each round, there is a $\frac{7}{8}$ probability that the supergame will continue for an additional round.[14] 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.[15]

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 in the experiment. 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.[16] 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 action choice of the matched subject.[17] The belief question is presented

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 $-1.41 = (22 - 39)/(51 - 39)$.

[13]The parameters used in this paper are identical to those used in one treatment of Embrey et al. [2018].

[14]The expected length of a supergame is hence eight rounds. The random termination is determined by a pseudo-random number generator whose seed is set arbitrarily at the beginning of the session.

[15]This method was first introduced in Fréchette and Yuksel [2017] and has now been used in multiple papers on a variety of topics. As in Vespa and Wilson [2019], we only use the method for the first *block*.

[16]Recall that μ_i^t is the probability assigned by i to j's choice of action C in round t.

[17]Unlike the classical quadratic scoring rule that is incentive compatible only under risk neutrality, incentive compatibility of the binarized scoring rule is independent of a subject's risk attitude. See

on a separate screen after subjects have made their action decision for that round and before feedback is provided.[18] 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.[19]

We conducted eight sessions per treatment and 16 sessions in total.[20] Table 2 summarizes basic information about each session. The supergames for the part with belief elicitation are separated into *early* and *late*. We use this categorization in the presentation of results, with most of the data analysis focusing on late supergames.[21] 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.[22]

Hossain and Okui [2013] and Allen [1987] for an earlier formulation of the idea. We use the implementation outlined in Wilson and Vespa [2018].

[18]We opted for this ordering where beliefs are elicited *after* action decisions in each round to minimize the risk that the belief questions influence the way subjects play these games.

[19]The maximum possible earning from this task is 99 experimental currency units.

[20]This number of sessions per treatment is more than the typical number. The reason for having more sessions will become apparent in the section on beliefs over strategies, because the method we propose is data intensive.

[21]We aimed for three supergames for both early and late when possible. When that was not possible, we aimed for each group to have a division of total rounds that was as balanced as possible.

[22]To address hedging concerns, we chose the supergame for the belief-elicitation task from the supergames not used for the action task. Experimental currency units were translated into earning in dollars at an exchange rate of 3 cents per point. 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 C) 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].

Table 2: Session Summary

Treatment	Session	No. of Subjects	No. of Supergames	No. of Game Rounds: Actions Only	No. of Game Rounds: Actions and Beliefs: Early	No. of Game Rounds: Actions and Beliefs	No. of Game Rounds: Actions and Beliefs: Late	Total no. of Obs. Rounds
Finite	1	20	12	8, 8, 8, 8	*8, 8, 8*	8, 8,	8, 8, 8	96
	2	20	12			8, 8,		96
	3	20	13			8, 8, 8,		104
	4	20	11			8,		88
	5	20	13			8, 8, 8,		104
	6	20	13			8, 8, 8,		104
	7	20	12			8, 8, 8,		104
	8	18	12			8, 8,		96
Indefinite	1	20	10	9, 7, 13, 7	*1, 2, 23,*	4, 14,	4, 1, 19	112
	2	20	9	8, 15, 7, 32	*2, 10,*		5, 1, 8	105
	3	18	7	8, 2, 3, 14	*25,*		17, 10	90
	4	16	8	9, 7, 10, 13	*32,*		7, 7, 6	96
	5	14	12	7, 22, 7, 3	*2, 5, 8,*		9, 3, 10	119
	6	14	6	1, 31, 4, 3	*24,*		15	94
	7	18	10	5, 6, 7, 14	*30, 8, 5,*		4, 9, 4	109
	8	20	9	11, 1, 4, 13	*9, 5,*		2, 4, 2	81

302 subjects in total.
Payment: \$8 + choices from two supergames (pre/post) + beliefs in one.
Earnings from \$22.00 to \$63.75 (with an average of \$35.30).
How to read: In the Finite treatment, session 1 had 20 subjects, they played a total of 12 supergames: 4 supergames of 8 rounds without belief elicitation, in the remaining 8 supergames that follow and where beliefs are) are also elicited, the first three (each with 8 rounds) are labelled "Early" supergames, three (each with 8) rounds are labelled "Late" supergames and the two in between (supergames 8 and 9—each having 8 rounds) fall in neither Early or Late category. In total subjects in that treatment played 96 rounds.

4 Results

The analysis of our data is separated into three sections. Section 4.1 provides a brief overview of the qualitative features of observed behavior focusing on actions. Section 4.2 presents results on beliefs (over actions), namely, their accuracy, how they are affected by history, and their relation to actions. Section 4.3 proposes a methodology to recover beliefs over supergame strategies and uses this method to study how the strategy choice relates to beliefs.

4.1 Actions

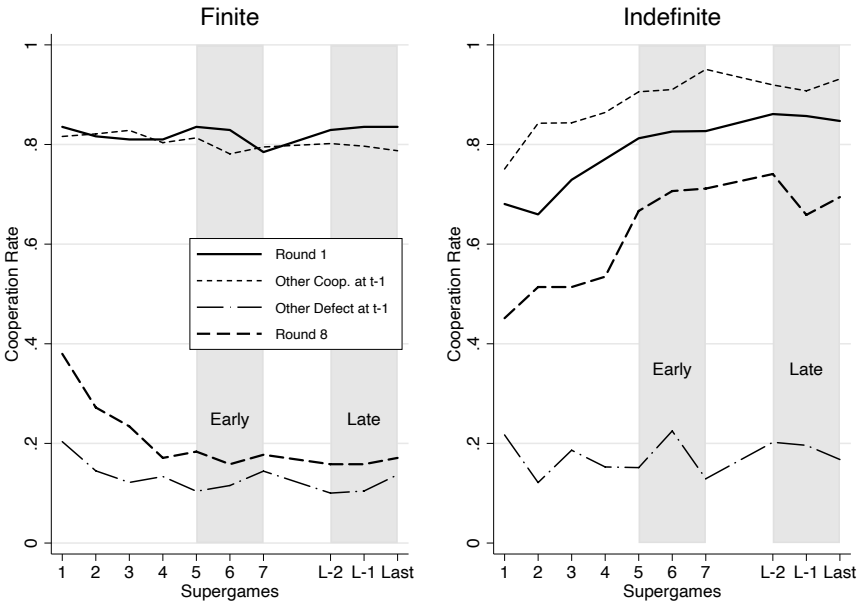

Figure 1: Cooperation Rate over Supergames

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 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.[23]

Hence, 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. Finally, cooperation collapses at the end of the Finite game but persists in the Indefinite game. In summary, behavior along key dimensions is qualitatively consistent with prior findings on these two games.[24]

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.

[23]Instead of round eight, one might want to compare the round-eight behavior in the Finite game to the last game round in the Indefinite game, or to the last observation round. Doing so does not qualitatively change the results. These alternative figures are presented in Online Appendix B (Figure 12).

[24]As mentioned in the Design section, Table 10 in the Online Appendix B also shows no significant changes in round-one choices for supergames where beliefs are elicited (compared with those where they are not) and Figure 13 shows that cooperation rates in the Finite game of when beliefs are elicited are not different from those in a prior experiment [Embrey et al., 2018] with the same parameters but without belief elicitation. Other key features of behavior in this environment, such as the rise in the popularity of threshold strategies with experience, are also observed in this dataset.

4.2 Actions and Beliefs

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

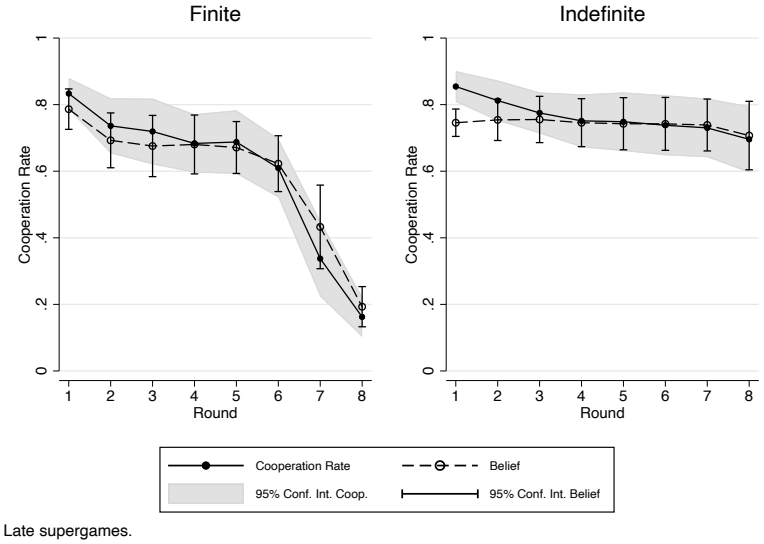

Figure 2: Choices and Beliefs by Round

Putting beliefs and actions together reveals beliefs–on average–track cooperation rates closely.[25] 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 rounds, the differences between action frequencies and beliefs are small, at less than one percentage point for Finite and two percentage points over the first eight rounds of the Indefinite game.[26] This difference is not statistically different from 0 for the Finite game, but it is for the Indefinite game (even though the difference is small in

[25]The reader interested in an equivalent to Figure 1 from the previous section but focusing on beliefs instead of actions is referred to Online Appendix B (Figure 14).

[26]Throughout, results over rounds will focus on the first eight rounds. For the Indefinite game, we have many more rounds, but sample sizes are substantially smaller for rounds nine and above.

magnitude).[27,28]

However, when we look at each round separately, both in the Finite and Indefinite 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 0.2 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. (The corresponding values are 0.1 and 1.1 in the Indefinite game.) 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 early rounds for the Indefinite game.

One natural question is whether, with experience, subjects learn to correct their mispredictions. Figure 3 displays the error in key rounds for early versus late supergames. As the figure shows, in many cases where more substantial error occurs in early supergames, improvement is observed in late supergames, but not for round seven of the Finite game and round one of the Indefinite game. Even in these cases, however, subjects' beliefs do move in the right direction. As seen in Figure 15 in Online Appendix B, which reports average cooperation rates and average beliefs for rounds one and seven over supergames, beliefs move in the correct direction with experience, but not fast enough to catch up with the changes in actions. We should note, however, that the changing behavior over the course of the session does not always imply beliefs are systematically off. For instance, in that same figure, one can see cooperation rates in round seven of the Indefinite game are changing with experience, but subjects correctly anticipate this change, as reflected in their beliefs.

Although determining exactly how beliefs are formed is not the goal of this study, understanding what allows subjects to predict actions relatively well is of clear in-

[27]We 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.

[28]Throughout, 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 Fréchette [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.

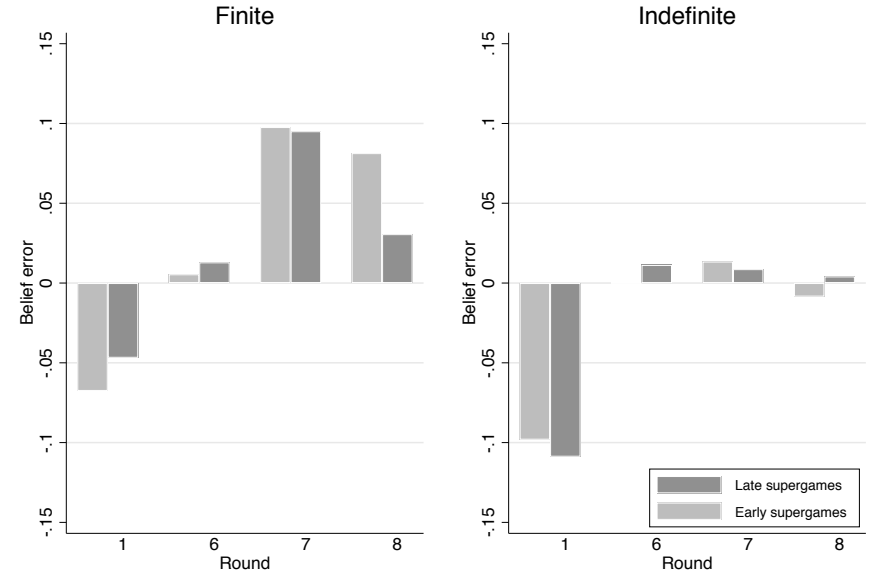

Figure 3: Belief Errors in Early vs. Late Supergames

terest. One conjecture is that subjects are simply reporting back their observations about others' behavior from previous supergames. Alternatively, subjects may form beliefs relying on introspection alone, or some combination of learning and introspection.[29] The data suggests that although experiences matter in shaping beliefs, they are not the sole determinant. Figure 16 in Online Appendix B shows the kernel density estimates of the differences between beliefs and the subject-specific experienced frequencies for the fifth (the first with belief elicitation) and last supergames of any given session. The figure reveals that subjects' beliefs differ substantially from the cooperation rates they have experienced. Consider round one where learning from past experiences is easiest (because there is no need to condition on history). In that round, beliefs differ from experienced frequencies by 17 and 16 percentage points, respectively in the first and last supergames (with belief elicitation) of the Finite game and 21 and 20 percentage points in the Indefinite game. This means that in many cases (58 percent) beliefs are further than plus or minus 20 percentage points

[29]The earlier observation about the Finite game—although behavior is changing in round seven, beliefs track action frequencies closely—already suggests subjects cannot be basing their beliefs only on empirical frequencies.

of the experienced cooperation rates.

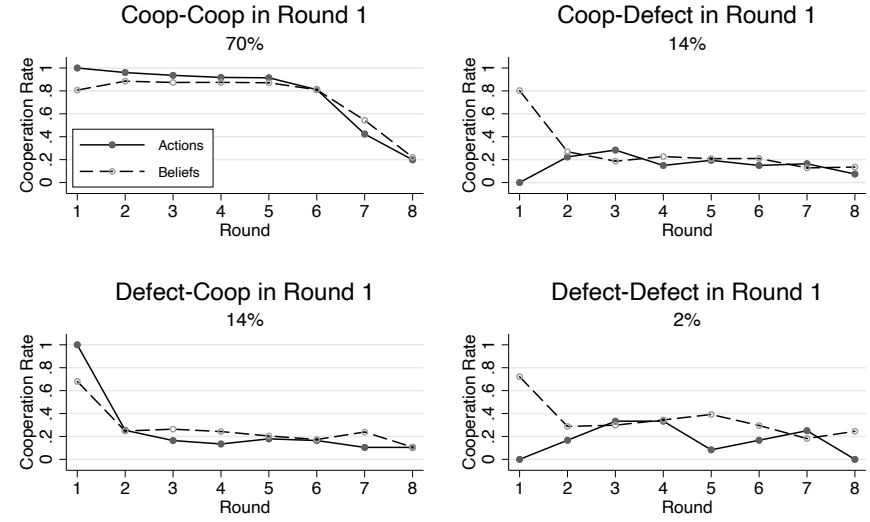

Own choice listed first.
Late supergames.
Percentage of cases under the title.
By construction the 'Actions' rate is 1 in round 1 of the left panels and 0 in the right panels.

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

So far in Figures 2 and 3, 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 question for round two. Figures 4 and 5 present the relevant data conditional on round-one histories (labeled with one's own action first followed by the opponent's action). In both the Finite and Indefinite games, we observe that beliefs quickly adjust in response to the other's action.[30] 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 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

[30]They should not be correct in round one, because beliefs are unconditional, whereas by construction, the figures present specific action frequencies in round one.

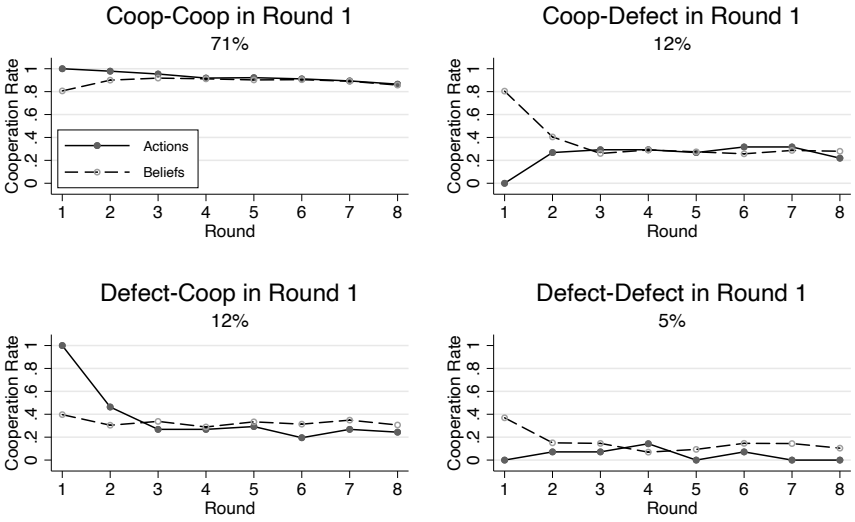

Own choice listed first.
Late supergames.
Percentage of cases under the title.
By construction the 'Actions' rate is 1 in round 1 of the left panels and 0 in the right panels.

Figure 5: Beliefs Conditional on Round One Action Pair, Indefinite Games

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 second round.

These results showing beliefs that are fairly accurate, both averaged over histories and along specific histories, do not speak directly to whether many or few subjects correctly anticipate actions at the individual level. One way to answer this question in a simple but structured way is to look at whether subjects are accurate in at least assessing whether cooperation by their opponent is a relatively likely or unlikely event. Specifically, we denote cooperation (by one's opponent) conditional on a history to be *unlikely* if the empirical frequency of cooperation is less than one third, *likely* if the empirical frequency is more than two thirds, and *uncertain* if the empirical frequency is between these values. Then, we identify the share of observations for which a subject's belief is accurate relative to this categorization; that is, we look at whether the belief lies in the same tercile (unlikely/likely/uncertain) as

the observed average cooperation rate. We do so for rounds one and two.

Table 12 in Online Appendix B shows that accuracy of beliefs at the individual level, as defined above, is high both for round one (73% in the Finite game, 67% in the Indefinite game) and round two (83% in the Finite game, 80% in the Indefinite game). The accuracy rate is substantially above 33% (the benchmark if beliefs were random) and this is true even in early supergames (above 65% in rounds 1 and 2 for both treatments). However, after one history, accuracy is low: in round two of the Indefinite game along $h^1 = (C, D)$ (cooperation by oneself and defection by the other), beliefs fall in the correct tercile only 29% of the time. Interestingly, the opposite is not true: round-two beliefs along $h^1 = (D, C)$ (defection by oneself and cooperation by the other) fall in the correct tercile 79% of the time. Table 12 also considers more demanding tests of accuracy by reporting the fraction of times the empirical frequencies of cooperation are within ± 5 and 10 percentage points of reported beliefs. Beliefs are fairly accurate along some histories (especially the more common ones, e.g., $h^1 = (C, C)$), but less so along other histories that are less common (particularly along $h^1 = (C, D)$ and (D, C) in the Indefinite game).

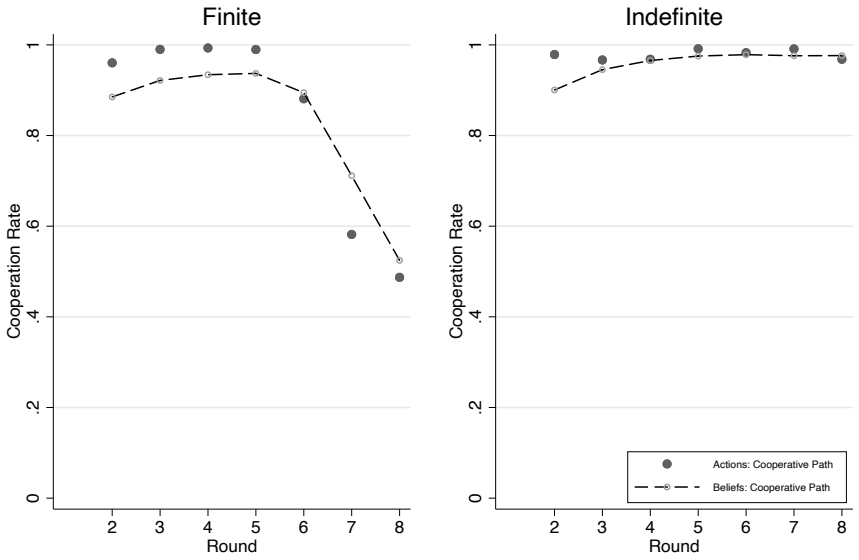

Cooperative path: joint cooperation up to that round.
Late supergames.

Figure 6: Cooperative Path (First Eight Rounds)

As Figures 4 and 5 above show, supergames starting with joint cooperation are the most common. How do beliefs evolve on a mutual cooperation path? Figure 6 shows the average cooperation rates $\bar{x}^t$ and average beliefs $\bar{\mu}^t$ along the history $h^{t-1} = ((C,C),\ldots,(C,C))$.[31] 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.[32] 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 points of the empirical frequency), they show optimism (13 percentage points higher than the empirical frequency) in round 7.[33] 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 supergames unfold. We note also these patterns are already visible in early supergames (see Figure 17 in Online Appendix B).

The last observation suggests, in particular, that the evolution of beliefs in the Finite game cannot simply be explained by heuristic models based on past action choices (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 and less than half as accurate than the ones we observe in the data.[34] 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 optimistic in later rounds of the Finite game and 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,

[31] Note $t = 2$ corresponds to cases presented in Figures 4 and 5.

[32] The 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.

[33] By round 8, the error declines to less than 4 percentage points.

[34] For the first exercise, we compare $\frac{1}{n}\sum_i (\mu_i^7 - x_{m(i)}^7)$ with $\frac{1}{n}\sum_i (x_{m(i)}^6 - x_{m(i)}^7)$, where $m(i)$ is the subject matched with subject i. For the second exercise, we compare $\frac{1}{n}(\sum_i |\mu_i^7 - x_{m(i)}^7|$ with $\frac{1}{n}\sum_i |x_{m(i)}^6 - x_{m(i)}^7|$.

subjects correctly anticipate cooperation will break down despite a history of joint cooperation in the Finite game.

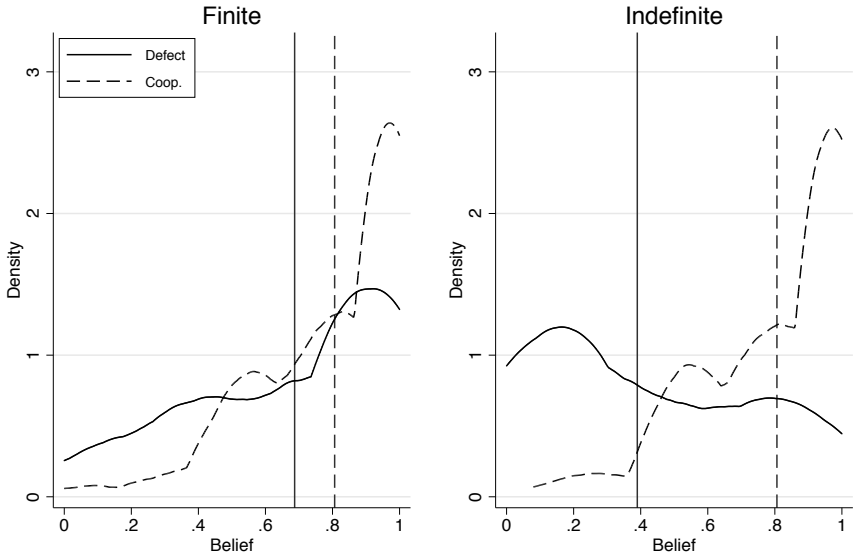

Figure 7: Beliefs of Defectors vs. Cooperators in Round One

We now turn to the question of whether different actions are supported by different beliefs. Figure 7 shows kernel density estimates of the distribution of round-one beliefs μ_i^1 by treatment and by the subject's own action a_i^1 in round one. At a broad level, we can easily see that the beliefs of cooperators and defectors are more different from one another in the Indefinite game than in the Finite game. The average beliefs of cooperators and defectors are statistically different in the Indefinite game ($p < 0.01$) but not in the Finite game.[35] Of those subjects who reported a belief of less than 50% in round one, only 39% cooperated in the Indefinite game, in contrast to 55% who cooperated in the Finite game. Hence, many subjects do not report beliefs that correspond to their own action. Subjects with optimistic beliefs cooperated in both treatments: of those subjects who reported a belief greater than 50% in round one, 94% cooperated in the Indefinite game and 87% cooperated in the Finite

[35]However, a Kolmogorov-Smirnov test rejects that the distributions are the same in both treatments at the 1% level, but we cannot account for the panel structure of the data with this test.

game. In other words, round-one beliefs were more predictive of round-one actions in the Indefinite game than in the Finite game, and subjects with higher beliefs tend to defect more often in the Finite game than in the Indefinite game.[36]

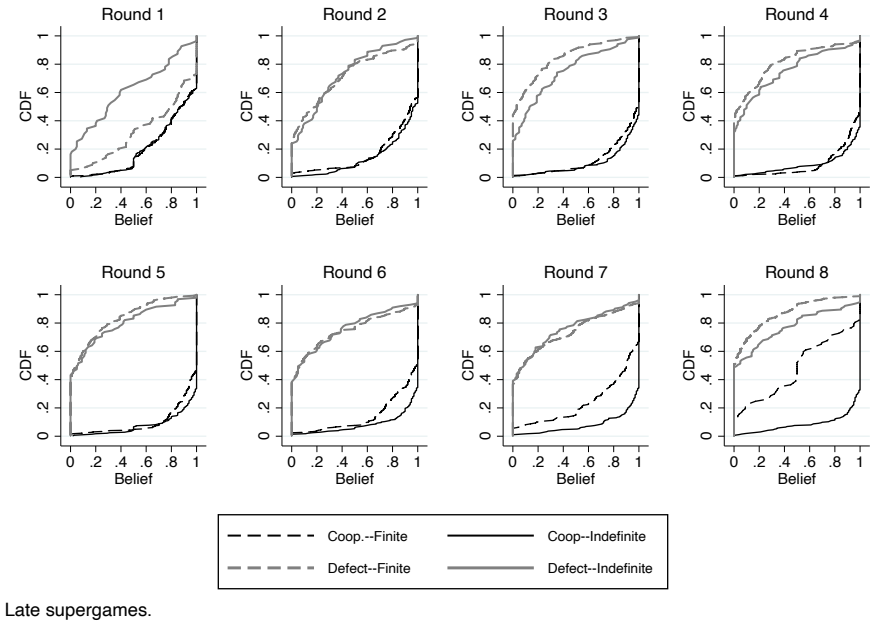

Figure 8: Beliefs by Action and Treatment: Rounds One through Eight

Figure 8 plots the CDF of beliefs by action and treatment for each round. It clearly shows cooperation and defection are associated with different beliefs. Except for round one in the Finite game, in every other comparison—every round for each treatment—the average belief is statistically different between those who cooperate versus those who defect (all p-values < 0.01). Higher cooperation rates are associated with more optimistic beliefs more generally. Table 13 in Online Appendix B shows that in all rounds, the marginal impacts of beliefs on the likelihood of cooperation are

[36]Appendix B reports additional results with respect to the determinants of cooperation using regression analysis. These results support the patterns reported in the paper. Specifically, we find beliefs are predictive of actions in both the Finite and Indefinite games. Focusing on round one, although beliefs are significant in both games, we find they have more predictive power in the Indefinite game. These results also suggest risk preferences have some limited predictive power for round-one choice in the Finite game (with the likelihood of cooperation decreasing with risk aversion). Recently, Proto et al. [2019] provide some evidence consistent with this finding.

positive in both the Finite and Indefinite games. In addition, the round number has a significant negative impact on cooperation in the Finite game but is insignificant in the Indefinite game.[37] More specifically, subjects are more likely to defect later in the Finite game, even if their beliefs are the same.

Perhaps more revealing, cooperation and defection in certain rounds are associated with different beliefs for Finite versus Indefinite games. In round eight, beliefs of the subjects who cooperate are statistically different across treatments ($p < 0.01$), as are those of the subjects who defect ($p < 0.1$). Subjects who defect in round eight of the Finite game are more pessimistic (on average) than those who do so in the Indefinite game. Similarly, subjects who cooperate are more optimistic in the Indefinite game than those in the Finite game. On the other hand, subjects who defect in round one of the Finite game are more optimistic than those who do so in the Indefinite game ($p < 0.01$). Hence, the same action can be supported by different beliefs in those two games.

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.

4.3 Beliefs over Supergame Strategies

The preceding section finds a link between beliefs and actions, and also that beliefs are not just the summary of past action choices in a supergame. These patterns lead us to the consideration of beliefs over strategies. Note that we do not make the claim that subjects reason in terms of strategies *per se*, but that we can represent their behavior as such.

The estimation method we develop has two stages and treats separately data on actions and beliefs without imposing any structure between them, thus allowing meaningful questions about whether strategies best respond to beliefs.

Plan for the Estimation Strategy:

1. Classify subjects into types on the basis of their choices.
   (a) Estimate strategies at the population level.
   (b) Use these estimates and each subject's choices to classify them.
2. Estimate beliefs over supergame strategies separately for each type.

[37] Specifications with our measure of risk attitude do not find that it is statistically significant.

Details are provided in sections 4.3.1, 4.3.2, and 4.3.3 below. Note that step 1 above can be performed in different ways. The version specified here as (a) and (b) are specific versions, but the method is more general than this. An example of this is given in Online Appendix B.4 where an alternative classification into types is performed. This exercise demonstrates the robustness of our results: while the specific method of typing that is adopted is entirely different, the qualitative results discussed below remain.

Here, we outline the intuition for the approach 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 that strategy is the one considered that 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.[38] How do we specify

[38]Note 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. Traditionally, in

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. We use results from this literature to determine which strategies to include in our 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.[39]

4.3.1 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. 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. We identify the values of p and β that maximize the likelihood of the observed sequences of action choices.

We use a two-step procedure to determine the set of strategies in our analysis. First we rely on prior evidence to construct a consideration set of 16 strategies. The consideration set includes all strategies that Fudenberg et al. [2012] report have a statistically significant SFEM estimate in at least one indefinitely repeated game with

the literature studying strategies in repeated games, such an assumption is introduced either at the design stage by directly restricting the set of strategies available to the 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.

[39]To validate this assumption, we do the following exercise. We compute the spread of beliefs defined as the difference between the 25^{th} and 75^{th} 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 is not in line with the assumption that the spread in beliefs is less among subjects of the same type.

perfect monitoring.[40] Motivated by the results of Embrey et al. [2018], who document the prevalent use of threshold strategies with experience in finitely repeated PD games, we also add to the consideration set all threshold strategies up to T8.[41] Results on this consideration set are reported in Online Appendix B. However, because our primary goal is to estimate beliefs over strategies, focusing on such a large set is more costly than is typical with SFEM: having more strategies can make identifying beliefs over different strategies difficult; it can also reduce the number of observations per type in the belief estimation. For these reasons, we use results from the larger consideration set to focus our analysis on the 10 strategies that are most important in terms of choices as well as beliefs. This set consists of AD, AC, Grim, TFT, STFT, Grim2, and TF2T, as well as threshold strategies T8, T7, and T6.[42]

Table 3: Strategy Prevalence and Typing

Finite			Indefinite		
	Share			Share	
Type	SFEM	Typing	Type	SFEM	Typing
T7	0.30	0.35	TFT	0.36	0.59
T8	0.22	0.20	Grim	0.18	0.09
AD	0.12	0.12	Grim2	0.11	0.11
TFT	0.09	0.12	AC	0.11	0.05
T6	0.08	0.08	TF2T	0.10	0.01
Grim	0.08	0.02	AD	0.09	0.10
TF2T	0.04	0.04	STFT	0.04	0.04
STFT	0.03	0.03	T8	0.01	0.01
AC	0.03	0.03	T7	0.00	0.00
Grim2	0.00	0.01	T6	0.00	0.00

Estimation using late supergames.
SFEM estimate for β are 0.94 for both.

[40] Our aim was to be inclusive in the first step of the selection process. In particular, our selection criterion is such that we include all the strategies found to be important in a variety of different papers that have estimated strategies and covered in the meta-study of Dal Bó and Fréchette [2018]. It also means that we do not include strategies that are not observed in direct elicitation studies (Dal Bó and Fréchette [2019] and Romero and Rosokha [2019]).

[41] Thus, the consideration set is AD, AC, Grim, TFT, STFT, Grim2, Grim3, TF2T, 2TFT, and T2–T8. Appendix B provides a detailed description of each of these strategies.

[42] From the original set, we eliminate T2–T5, which our estimates indicate are not relevant in the Finite game, as well as 2TFT and Grim3, which are not popular enough in the Indefinite game to generate reliable belief estimates.

Table 3 presents the estimation results (in columns 2 and 5) sorted by prevalence. The 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].[43]

More specifically, 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.[44] 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).

Result 4 We reproduce results about strategy choices observed in previous finitely and indefinitely repeated PD games.[45] 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.

4.3.2 Typing of Subjects

We use the SFEM results to type subjects according to the strategy that they are most likely playing.[46] Recall the SFEM yields the probability distribution over su-

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

[44]We use *conditionally cooperative* as distinct from strategies that unilaterally defect or strategies that do not condition on the choices of the opponent.

[45]To our knowledge, this study is the first to compare strategies in Finite and Indefinite games within the same experimental paradigm.

[46]Dvorak [2020] recently provides a R-package for easy implementation of the SFEM using the EM algorithm, which also includes a similar typing procedure.

pergame strategies (p) and the probability of implementation errors ($1-\beta$). These probabilities can be used to compute the Bayesian posterior that a subject is playing each of the candidate supergame strategies given the sequence of his actions. Each subject is associated with the supergame strategy that has the highest likelihood according to this posterior.[47]

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 are then given by $\widehat{AD}$ and $\widehat{AC}$, where for every h^{t-1},

$$\widehat{AD}(h^{t-1}) = 0.9 \circ D + 0.1 \circ C,$$
$$\widehat{AC}(h^{t-1}) = 0.9 \circ C + 0.1 \circ D.$$

We suppose the strategy of each subject is chosen from the set $\widehat{Z} = \{\widehat{AD}, \widehat{AC}\}$ using the prior distribution p. Assume now that a subject exists 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, the posterior that the subject is playing $\widehat{AD}$ is $\frac{p_{\widehat{AD}}\beta^4(1-\beta)^{20}}{p_{\widehat{AD}}\beta^4(1-\beta)^{20}+p_{\widehat{AC}}\beta^{20}(1-\beta)^4}$, 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. Note that in the actual typing exercise, most of the strategies are history dependent. This finding implies that calculating the Bayesian posterior requires comparing for each round the actual action choice of the subject with the action implied by each strategy given the history up to that point.

The results of the typing exercise are reported in the third and sixth columns of Table 3. 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.[48] Clearly, the smaller the fraction of subjects of a given type, the less reliable their belief estimates will be.

[47] A unique strategy exists within the consideration set for each subject in our data set that achieves the highest posterior (given the SFEM results). In Online Appendix B, we plot the CDF of the posterior given the type of the subject. We also plot the CDFs of the second most likely type, third, etc., illustrating the ability of the data to identify types.

[48] Two potential sources for such differences are possible. First, and simply mechanically, some subjects play more supergames than others; thus, 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 play TFT, and a small fraction plays Grim. However, for

4.3.3 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.[49] 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 $1-\tilde{\beta}$. For every h^{t-1},

$$\tilde{\sigma}_j^t(h^{t-1}) = \begin{cases} (\tilde{\beta}) \circ C + (1-\tilde{\beta}) \circ D & \text{if } \sigma_j^t(h^{t-1}) = C, \\ (\tilde{\beta}) \circ D + (1-\tilde{\beta}) \circ C & \text{if } \sigma_j^t(h^{t-1}) = D. \end{cases}$$

Note $\tilde{p}$ 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 $\widetilde{\text{AD}}$ and $\widetilde{\text{AC}}$ for $\tilde{\beta} = 0.9$. It then follows that the round-one belief μ_i^1 equals $\tilde{p}_{\widetilde{\text{AD}}} \times 0.1 + \tilde{p}_{\widetilde{\text{AC}}} \times 0.9$. If the subject observes $a_j^1 = C$ in the first round, by Bayes' rule, his belief in round two will increase to

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

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. A summary of these estimation results are reported in Tables 4 and 5, with the complete results provided in the Appendix. Note 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

some of the subjects playing Grim, the number of observations that distinguishes Grim from TFT is very small. When computing the posterior at the subject level, the few observations of difference for a given subject may not be enough to generate the highest posterior on Grim given the strong prior in favor of TFT.

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

reported in Tables 4 and 5. 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 subjects put the most weight on, on average).

Table 4: Beliefs over Strategies in the Finite Game

	Share		Estimated Beliefs - $\tilde{p}$									
Type	SFEM	Typing	T7	T8	Grim	TFT	AD	TF2T	Grim2	Other	ν	$\tilde{\beta}$
T7	0.30	0.35	0.43	0.39	0.18	0.00	0.00	0.00	0.00	0.00	0.03	1.00
T8	0.22	0.20	0.00	0.50	0.04	0.01	0.09	0.15	0.21	0.00	0.01	1.00
AD	0.12	0.12	0.75	*[0.00]*	*[0.00]*	0.00	0.07	0.00	0.00	0.18	0.03	1.00
TFT	0.09	0.12	0.00	0.33	0.00	0.53	0.11	0.00	0.00	0.03	0.01	1.00
T6	0.08	0.08	0.99	0.00	*[0.00]*	0.00	0.00	0.00	0.00	0.00	0.01	1.00
Grim	0.08	0.02	0.00	0.22	0.17	0.16	0.34	0.01	0.00	0.10	0.02	1.00
Other	0.11	0.11	0.01	0.16	0.35	0.30	0.01	0.11	0.00	0.06		
All			0.30	0.29	0.11	0.09	0.07	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 6.

Table 5: Beliefs over Strategies in the Indefinite Game

	Share		Estimated Beliefs - $\tilde{p}$									
Type	SFEM	Typing	Grim	TFT	TF2T	AC	AD	Grim2	STFT	Other	ν	$\tilde{\beta}$
TFT	0.36	0.59	0.28	0.25	0.19	0.00	0.08	0.14	0.05	0.00	0.02	1.00
Grim	0.18	0.09	0.80	0.13	0.02	0.00	0.00	0.05	0.00	0.00	0.02	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	0.03	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	0.05	1.00
Other	0.05	0.05	0.35	0.00	0.00	0.00	0.48	0.16	0.00	0.00		
All			0.32	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, andTF2T.
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 7.

Tables 4 and 5 reveal important differences in beliefs between the Finite and the 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, we observe subjects' beliefs to be in line with actual 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 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. These estimates are from late supergames; hence, our results indicate heterogeneity in beliefs across types can be persistent.

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 4 and 5 report detailed information on how beliefs differ by type in each game. However, the richness of these data makes it difficult to identify general patterns on how beliefs over strategies are connected to subjects' strategy choice. Figure 9 hones in on broad features of those beliefs. First, the left two panels display, for each of the most common types (ordered by popularity), the difference between their beliefs that others are of their type and the reality (as estimated by SFEM). As can be seen, subjects display a tendency to believe others are more like themselve than they actually are (the last bar indicates the weighted average). 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]. Note also, however, that this is not very pronounced (or even in the other direction) for the most common types: T7 for the finite game and TFT for the indefinite game.

Second, the other panels make use of a ranking of strategies 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 $\beta \to 1$), the expected payoff associated with playing the former strategy against itself is higher than the expected payoff of playing the later strategy against itself.[50] This generates the following order of cooperativeness (from least to most): AD, STFT, T6, T7,

[50] Analytical derivation of the cooperativeness order for an infinitely repeated PD with a general specification of the payoffs and discount factor is presented in Online Appendix D . We also

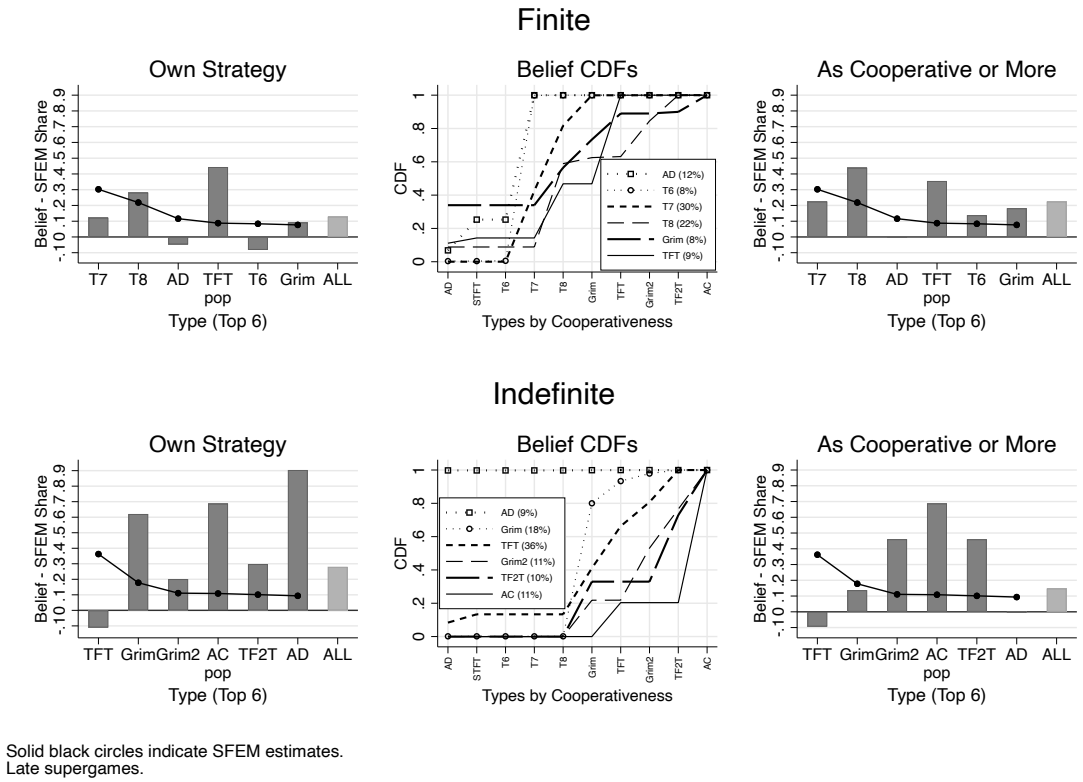

Solid black circles indicate SFEM estimates.
Late supergames.

Figure 9: Belief Patterns

T8, Grim, TFT, Grim2, TF2T, and AC. Note that for the strategies considered, the cooperativeness order is strict; as such, subjects' beliefs about others being as cooperative also correspond to their beliefs about others using the same strategy as them. The two central panels give the CDF of beliefs ranked by cooperativeness for each of the most popular types. Although not presenting exactly a first order stochastic dominance (FOSD) relation when going from less cooperative strategies to more cooperative ones, the general movement is in that direction. For instance, in the Finite game, AD, T6, and T7 types do show a FOSD relation, with more cooperative types being more optimistic about the cooperativeness of others. T8

describe ways to numerically verify that this order is preserved in the Finite game. On the subset of strategies considered by Proto et al. [2020], our cooperativeness order coincides with the inverse of their *harshness* ranking.

also broadly fits this structure and, albeit for a small overlap at the bottom, is to the right (more optimistic) of the distribution for T7. In the Indefinite game the comparisons are slightly more muddled, but some patterns emerge. For instance the beliefs of AC and TF2T types (the two most cooperative strategies) FOSD the beliefs of AD types. However, overall, the beliefs across the different cooperative strategies are not clearly ranked.

The right two panels also use the cooperativeness order. We compute for each type the belief that others are at least as cooperative as they are and compare it to the actual frequency (as estimated by SFEM). The panels report this difference for each of the most popular types. Most important types in both games underestimate the likelihood that others are less cooperative than they are. In particular, in the Finite game, the five cooperative types all believe others will not defect before they will. For certain types, this bias is large in magnitude.[51]

The accuracy of beliefs over strategies can be studied more directly without relying on the cooperativeness order. In Figure 28 of Online Appendix B, we compute, for each type, the Euclidean distance between beliefs and the estimated frequency of strategies. To study whether beliefs become more accurate with experience, we also look at how this distance changes from early to late supergames. We find that, in aggregate, beliefs are becoming more accurate with experience in the Finite game, whereas accuracy changes little in the Indefinite game. In both cases, the most popular strategy types (T7 in Finite and TFT in Indefinite) have the most accurate beliefs in late supergames.[52]

Additionally, in the same Online Appendix (Figures 21-27, Tables 19-20), to study learning effects more generally, we document in detail how the distribution of strategies, types, and beliefs for each type change from early to late supergames. We summarize the key observations from these results here. While behavior stabilizes quickly in the Indefinite game—with little change in distribution of strategies, types and beliefs observed from early to late supergames—there is clear evidence of learning in the Finite game. Most significantly, there is a shift towards less cooperative strategies: popularity of T8 declines while the popularity of T7 and T6 increase. The observed shift in strategies is anticipated by beliefs. In early supergames, the

[51]Note that in the case of AD, the fact that the difference is 0 is mechanical: all strategies are at least as cooperative as AD.

[52]In the Finite game, early beliefs overestimate the likelihood of T8 and underestimate the likelihood of T7. Both of these errors are reduced (or eliminated) with experience. For the Indefinite game, early beliefs overestimate the likelihood of Grim and underestimate the likelihood of TFT; however, these errors (which are less costly than those observed in the Finite game) are not corrected with experience.

aggregate belief weight on Grim and T8 are 21 and 49 percent, respectively. These weights decline to 11 and 29 percent, respectively in late supergames. By contrast, the aggregate weight on T7 increases from 6 percent to 30 percent. These results suggest, in the Finite game, subjects to be updating their beliefs about the cooperativeness of their counterpart throughout the session and adjusting their strategy choices in response to these changing beliefs.

Result 6 Substantial heterogeneity exists in beliefs within each game: subjects using different strategies hold different beliefs. The results also suggest subjects tend to overestimate the likelihood that others use the same strategy as their own, while underestimating the likelihood that others use less cooperative strategies.

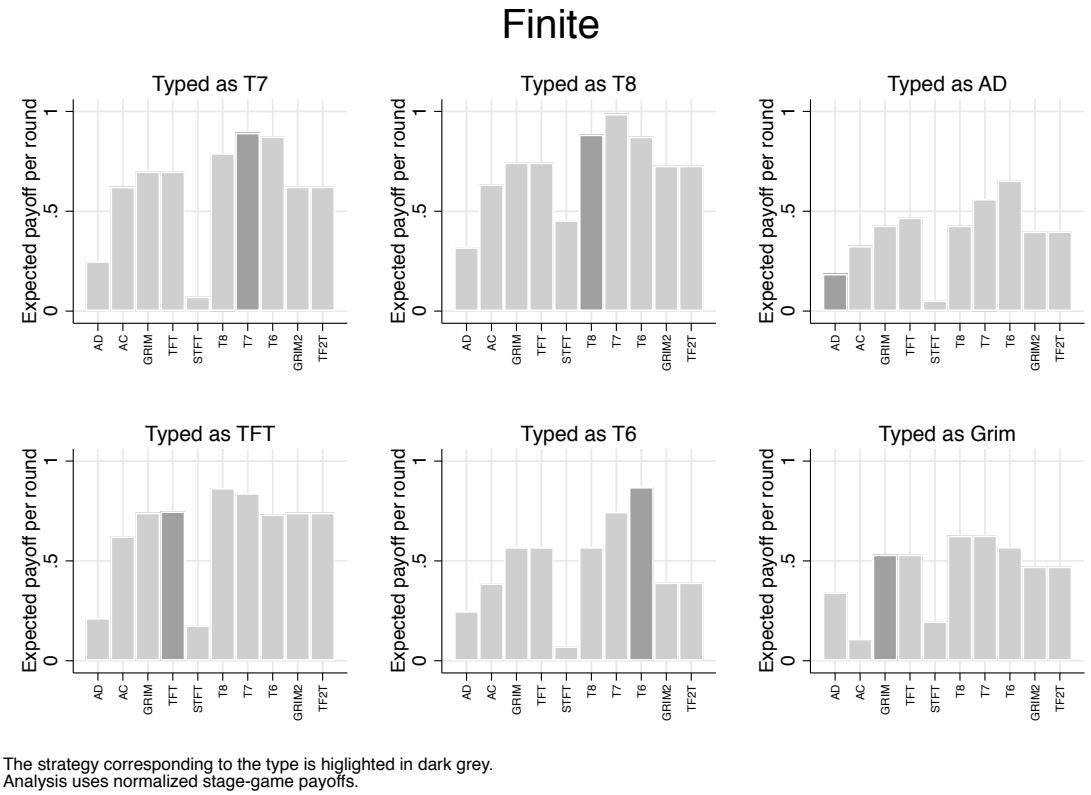

Figure 10: Best Response for Top 6 Types in the Finite Game

The observation that subjects using different strategies hold different beliefs raises the question of how they are connected. To shed light on this connection, we explore

the extent to which subjects are subjectively rational. That is, we study how close a subject's strategy is to being optimal within the set of supergame strategies we consider. 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. For the purposes of our discussion in this section, we consider *subjective rationality* in the constrained sense and examine if her strategy choice is a best response to her supergame beliefs *within* the set of strategies Z in the consideration set.[53]

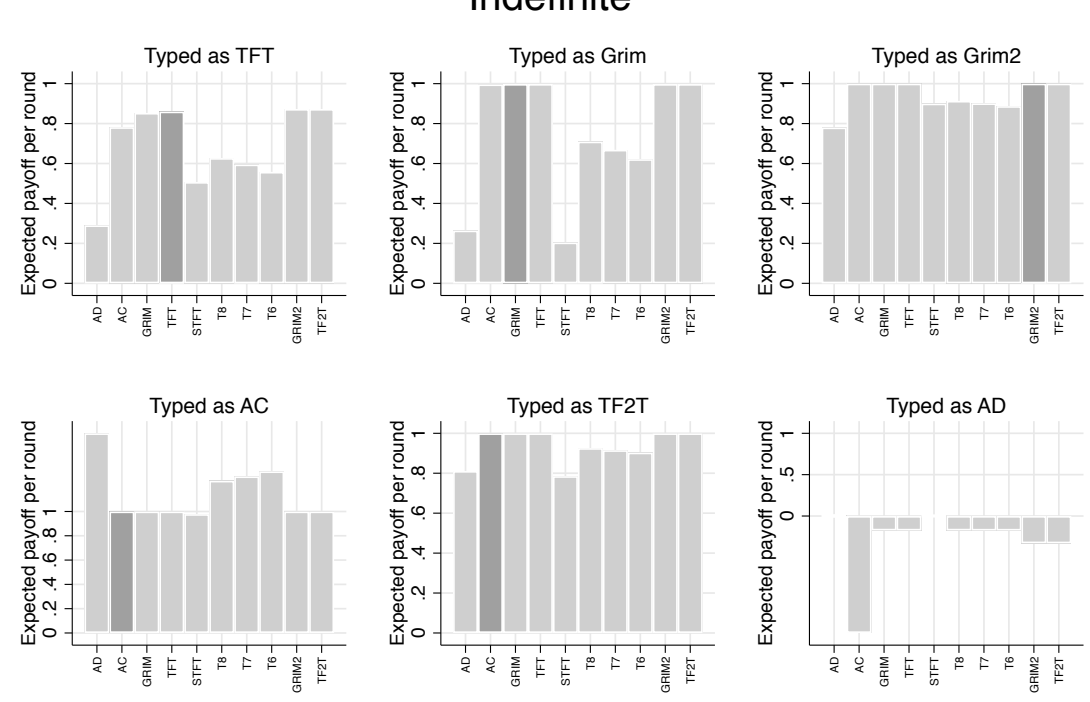

Figure 11: Best Response for Top 6 Types in the Indefinite Game

The results, presented in Figures 10 and 11, suggest most subjects' strategy

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

choices are either exact or approximate best responses given their supergame beliefs.[54] 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.[55]

In the Indefinite game, a similar pattern emerges. Most common types (TFT, Grim, Grim2, TF2T, and AD—84% of subjects) almost exactly best respond to their beliefs.[56] 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.[57] Hence, overall, the majority of subjects appear subjectively rational or close to subjectively rational.

Result 7 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 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

[54]Table 18 in the Online Appendix provides detailed best-response analysis for each of the six common types in both the Finite and Indefinite games.

[55]Note 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.

[56]For TFT, the strict best response is TF2T or Grim2, but TFT achieves 99% of the best-response payoff.

[57]The other type for which strategy choice is far from best response is STFT (4%). Given beliefs, the best response is TFT.

game only achieves 64% of the best-response payoff.

5 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 and beliefs over strategies. Our first key finding is that round beliefs are, in aggregate, remarkably accurate. In both the Finite and Indefinite games, beliefs averaged over all rounds are less than three percentage points away from the empirical action frequencies. Beliefs also adjust appropriately to the history of play even when these adjustments are not small: in some histories, they move by almost 60 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. Importantly, results show beliefs over stage actions are *forward looking*. Most notably, beliefs along the history of mutual cooperation evolve very differently in the Finite and the Indefinite games. Persistence of cooperation in the Indefinite game and its collapse in the Finite game are correctly anticipated along such histories. Interestingly, the same choice can be observed in both games in situations where subjects report very different beliefs.

Our second category of findings is based on the development of a novel method to recover beliefs over supergame strategies from beliefs over stage actions in each round. First, subjects in the Finite and Indefinite games correctly anticipate the different class of strategies used: threshold strategies in the former and conditionally cooperative strategies in the later. Second, subjects playing different strategies have strikingly heterogeneous beliefs over the strategy choice of the other player. This heterogeneity in strategies can be related to the heterogeneity in 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. These results also indicate that although, as noted above, beliefs are surprisingly accurate, since deviations tend to be systematic, they have specific implications. In particular, most types believe that others use strategies that are at least as cooperative as their own (or more). This points to one factor contributing to the slow unravelling of cooperation in the

Finite game. 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.

The new procedure proposed here to recover beliefs over strategies has broader applicability. In repeated games, it can be applied to different sets of strategies and/or be combined with alternative methods to type subjects. It does not require the use of the SFEM per se. More generally, this procedure can be used to recover beliefs over strategies from beliefs over actions in any sequential game. In simple enough games, the complete set of pure strategies can be included in the analysis. To demonstrate the versatility of the procedure in our current experiment, Online Appendix B.4 revisits the belief-estimation results under a simplified alternative approach. We take the most popular defective strategy (AD) and the most popular cooperative strategies (T7 for the Finite game and TFT for the Indefinite game) and classify a subject as being AD, T7, or TFT, if 90% of their choices across all supergames correspond to what the strategy dictates.[58] The results on estimated beliefs for these subjects presented in Figure 19 echo those in our main analysis. First, subjects in the Finite game expect others to play threshold strategies more than subjects in the Indefinite game. Second, beliefs display heterogeneity across types. Third, in the Finite game, T7 is subjectively rational while AD is not; in the Indefinite game both TFT and AD are subjectively rational. Fourth, in the Finite game TFT is the type that overestimates its own popularity the most while in the Indefinite game it is AD. Hence, the results presented in the paper are robust to an alternative methods to determine types.Finally, although the approach assumes Bayesian updating, we note that the results are not simply "noise" and in line with observations at the level of round beliefs. Although a strong assumption, it provides a useful starting point to discuss beliefs over strategies and has the substantial advantage of allowing fewer constraints and a less heavy handed intervention when eliciting beliefs.

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 for both Finite and Indefinite games, standard preferences with optimizing behavior but small errors in beliefs goes a long way to account for the observed behavior. This illustrates how small but systematic departures from accurate beliefs (at key points in the supergame) can sustain long-run cooperation in finitely repeated PD. Although beliefs are generally accurate, for 80% of subjects, best responding to their *subjective* beliefs (that are slightly over-optimistic) involves cooperating more than would be *objectively* optimal (given the observed strategy distribution in the population).

[58]The few subjects that could be classified as either T7 or TFT are dropped from the exercise.

Furthermore, in the indefinitely repeated PD, our results highlight the difficulty of resolving equilibrium selection. Different subjects hold persistently different beliefs about others, and in environments conducive to cooperation, such as ours, they experience few histories where those beliefs are revealed to be incorrect. As a consequence, a variety of conditionally cooperative strategies remain popular despite many repetitions.

In summary, our results on beliefs suggest subjects understand the difference between finitely and indefinitely repeated environments even when their observed behavior in terms of actions is identical. 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.

Appendix: Complete Estimation Results

Table 6: Estimates for the Finite Game on Late Supergames

	Share		Estimated Beliefs - $\tilde{p}$											
	SFEM	TYPING	AD	AC	GRIM	TFT	STFT	T8	T7	T6	GRIM2	TF2T	ν	$\tilde{\beta}$
T7	0.30	0.35	0.00	0.00	0.18	0.00	0.00	0.39	0.43	0.00	0.00	0.00	0.04	1.00
			(0.01)	(0)	(0.14)	(0.09)	(0)	(0.16)	(0.15)	(0)	(0)	(0.01)		
T8	0.22	0.20	0.09	0.00	0.04	0.01	0.00	0.50	0.00	0.00	0.21	0.15	0.04	1.00
			(0.08)	(0.06)	(0.1)	(0.14)	(0.02)	(0.12)	(0.07)	(0)	(0.11)	(0.11)		
AD	0.12	0.12	0.07	0.00	*[0.00]*	0.00	0.18	*[0.00]*	0.75	*[0.00]*	0.00	0.00	0.06	1.00
			(0.09)	(0.02)		(0.1)	(0.09)		(0.17)		(0.04)	(0.06)		
TFT	0.09	0.12	0.11	0.00	0.00	0.53	0.03	0.33	0.00	0.00	0.00	0.00	0.05	1.00
			(0.07)	(0.05)	(0.06)	(0.22)	(0.04)	(0.12)	(0.03)	(0)	(0.05)	(0.13)		
T6	0.08	0.08	0.00	0.00	*[0.00]*	0.00	0.00	0.00	0.99	0.00	0.00	0.00	0.03	1.00
			(0.07)	(0)		(0.21)	(0.02)	(0.15)	(0.31)	(0.09)	(0)	(0.01)		
GRIM	0.08	0.02	0.34	0.10	0.17	0.16	0.00	0.22	0.00	0.00	0.00	0.01	0.07	1.00
			(0.2)	(0.05)	(0.34)	(0.08)	(0.12)	(0.13)	(0.09)	(0.01)	(0.03)	(0.04)		
TF2T	0.04	0.04	0.00	0.14	0.83	0.00	0.00	0.00	0.03	0.00	0.00	0.00	0.04	1.00
			(0.09)	(0.1)	(0.35)	(0.1)	(0.01)	(0.13)	(0.11)	(0.03)	(0.17)	(0.07)		
STFT	0.03	0.03	0.00	0.00	*[0.00]*	0.65	0.00	*[0.00]*	0.00	0.00	0.00	0.35	0.11	1.00
			(0.02)	(0.13)		(0.4)	(0.02)		(0.03)	(0.02)	(0.07)	(0.38)		
AC	0.03	0.03	0.04	0.00	0.16	0.30	0.03	0.46	0.00	0.00	0.00	0.00	0.07	1.00
			(0.04)	(0.05)	(0.13)	(0.16)	(0.04)	(0.2)	(0.03)	(0.03)	(0.04)	(0.12)		
GRIM2	0.02	0.01	-	-	-	-	-	-	-	-	-	-	-	-
ALL			0.07	0.01	0.11	0.09	0.03	0.29	0.30	0.00	0.05	0.05		

Estimation on late supergames. SFEM estimate for β is 0.94.
Estimates in *[square brackets]* are not estimated due to collinearity.
Estimates in *(brackets)* show bootstrapped standard deviation.

Table 7: Estimates for the Indefinite Game on Late Supergames

	Share		Estimated Beliefs - $\tilde{p}$											
	SFEM	TYPING	AD	AC	GRIM	TFT	STFT	T8	T7	T6	GRIM2	TF2T	ν	$\tilde{\beta}$
TFT	0.36	0.59	0.08	0.00	0.28	0.25	0.05	0.00	0.00	0.00	0.14	0.19	0.01	1.00
			(0.04)	(0.06)	(0.14)	(0.13)	(0.03)	(0)	(0)	(0)	(0.12)	(0.12)		
GRIM	0.18	0.09	0.00	0.00	0.80	0.13	0.00	0.00	0.00	0.00	0.05	0.02	0.06	1.00
			(0.06)	(0.09)	(0.24)	(0.17)	(0.04)	(0.03)	(0.02)	(0.01)	(0.06)	(0.11)		
GRIM2	0.11	0.11	0.00	0.23	0.22	0.00	0.00	0.00	0.00	0.00	0.31	0.23	0.02	1.00
			(0.03)	(0.12)	(0.16)	(0.07)	(0.02)	(0)	(0)	(0)	(0.2)	(0.14)		
AC	0.11	0.05	0.00	0.80	0.00	0.20	0.00	0.00	0.00	0.00	0.00	0.00	0.11	1.00
			(0.04)	(0.33)	(0.03)	(0.26)	(0.05)	(0.02)	(0.02)	(0.02)	(0.14)	(0.11)		
TF2T	0.10	0.01	0.00	0.27	0.33	0.00	0.00	0.00	0.00	0.00	0.00	0.40	0.01	1.00
			(0)	(0.13)	(0.2)	(0.04)	(0)	(0.01)	(0)	(0)	(0.04)	(0.18)		
AD	0.09	0.10	1.00	0.00	0.00	0.00	0.00	*[0.00]*	*[0.00]*	*[0.00]*	0.00	0.00	0.04	1.00
			(0.24)	(0.02)	(0.1)	(0.05)	(0.17)				(0.01)	(0.01)		
STFT	0.04	0.04	0.48	0.00	0.35	0.00	0.00	0.00	0.00	0.00	0.16	0.00	0.08	1.00
			(0.27)	(0.07)	(0.24)	(0.12)	(0.11)	(0.05)	(0.04)	(0.06)	(0.09)	(0.09)		
T8	0.01	0.01	-	-	-	-	-	-	-	-	-	-	-	-
T7	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-
T6	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-
ALL			0.14	0.14	0.32	0.14	0.02	0.00	0.00	0.00	0.10	0.14		

Estimation on late supergames. SFEM estimate for β is 0.94.
Estimates in *[square brackets]* are not estimated due to collinearity.
Estimates in *(brackets)* show bootstrapped standard deviation.

Table 8: Estimates for the Finite Game on Early Supergames

	Share		Estimated Beliefs - $\tilde{p}$											
	SFEM	TYPING	AD	AC	GRIM	TFT	STFT	T8	T7	T6	GRIM2	TF2T	ν	$\tilde{\beta}$
T8	0.30	0.36	0.01	0.00	0.40	0.00	0.00	0.58	0.00	0.00	0.00	0.00	0.05	1.00
			(0.06)	(0.02)	(0.13)	(0.07)	(0.01)	(0.1)	(0.04)	(0.01)	(0.05)	(0.04)		
T7	0.25	0.20	0.00	0.00	*[0.00]*	0.00	0.00	0.75	0.25	0.00	0.00	0.00	0.03	1.00
			(0.03)	(0)		(0.11)	(0.02)	(0.17)	(0.12)	(0.01)	(0.01)	(0.02)		
TFT	0.17	0.15	0.19	0.00	0.50	0.00	0.03	0.26	0.00	0.00	0.02	0.00	0.05	1.00
			(0.09)	(0.02)	(0.21)	(0.15)	(0.05)	(0.14)	(0.04)	(0)	(0.05)	(0.07)		
AD	0.12	0.13	0.25	0.00	*[0.00]*	0.21	0.00	0.55	*[0.00]*	*[0.00]*	0.00	0.00	0.11	1.00
			(0.14)	(0.04)		(0.16)	(0.1)	(0.21)			(0.05)	(0.09)		
TF2T	0.05	0.08	0.15	0.00	0.06	0.54	0.15	0.11	0.00	0.00	0.00	0.00	0.05	1.00
			(0.19)	(0.05)	(0.12)	(0.19)	(0.16)	(0.09)	(0.01)	(0.02)	(0.04)	(0.06)		
GRIM2	0.04	0.03	0.00	0.30	0.00	0.43	0.00	0.27	0.00	0.00	0.00	0.00	0.02	1.00
			(0.02)	(0.19)	(0.2)	(0.18)	(0.04)	(0.19)	(0.02)	(0.03)	(0.1)	(0.1)		
STFT	0.03	0.04	0.00	0.00	*[0.00]*	0.42	0.00	0.00	*[0.00]*	*[0.00]*	0.58	0.00	0.15	1.00
			(0.09)	(0.12)		(0.33)	(0.12)	(0.1)			(0.33)	(0.18)		
AC	0.03	0.02	0.03	0.78	*[0.00]*	0.16	0.00	0.03	*[0.00]*	*[0.00]*	0.00	0.00	0.16	0.92
			(0.26)	(0.37)		(0.09)	(0.08)	(0.07)			(0.1)	(0.1)		
GRIM	0.01	0.00	-	-	-	-	-	-	-	-	-	-	-	-
T6	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-
ALL			0.07	0.04	0.21	0.09	0.01	0.49	0.06	0.00	0.02	0.00		

Estimation on early supergames. SFEM estimate for β is 0.92.
Estimates in *[square brackets]* are not estimated due to collinearity.
Estimates in *(brackets)* show bootstrapped standard deviation.

Table 9: Estimates for the Indefinite Game on Early Supergames

	Share		Estimated Beliefs - $\tilde{p}$											
	SFEM	TYPING	AD	AC	GRIM	TFT	STFT	T8	T7	T6	GRIM2	TF2T	ν	$\tilde{\beta}$
TFT	0.36	0.60	0.08	0.19	0.40	0.16	0.05	0.00	0.00	0.00	0.00	0.12	0.01	1.00
			(0.03)	(0.09)	(0.13)	(0.09)	(0.03)	(0)	(0)	(0)	(0.04)	(0.06)		
GRIM	0.21	0.09	0.11	0.11	0.45	0.19	0.14	0.00	0.00	0.00	0.00	0.00	0.10	1.00
			(0.11)	(0.2)	(0.22)	(0.14)	(0.09)	(0.04)	(0.11)	(0.08)	(0.12)	(0.16)		
TF2T	0.14	0.10	0.12	0.00	0.25	0.37	0.11	0.00	0.00	0.00	0.06	0.09	0.02	1.00
			(0.09)	(0.05)	(0.12)	(0.11)	(0.08)	(0.01)	(0.01)	(0.01)	(0.06)	(0.08)		
AD	0.13	0.13	0.59	0.03	0.20	0.00	0.14	*[0.00]*	*[0.00]*	*[0.00]*	0.04	0.00	0.05	1.00
			(0.22)	(0.03)	(0.11)	(0.05)	(0.14)				(0.04)	(0.05)		
GRIM2	0.10	0.05	0.42	0.00	0.31	0.00	0.00	0.00	0.00	0.00	0.26	0.00	0.06	1.00
			(0.23)	(0.13)	(0.2)	(0.05)	(0.07)	(0.05)	(0.02)	(0.02)	(0.23)	(0.07)		
AC	0.05	0.02	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.08	1.00
			(0.01)	(0.09)	(0.01)	(0.44)	(0.01)	(0.01)	(0.01)	(0.01)	(0.05)	(0.12)		
STFT	0.02	0.01	0.00	0.02	0.15	0.16	0.53	*[0.00]*	*[0.00]*	*[0.00]*	0.10	0.03	0.05	1.00
			(0.03)	(0.02)	(0.08)	(0.08)	(0.31)				(0.06)	(0.02)		
T8	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-
T7	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-
T6	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-
ALL			0.14	0.14	0.32	0.14	0.02	0.00	0.00	0.00	0.10	0.14		

Estimation on early supergames. SFEM estimate for β is 0.94.
Estimates in *[square brackets]* are not estimated due to collinearity.
Estimates in *(brackets)* show bootstrapped standard deviation.

References

Franklin Allen. Discovering personal probabilities when utility functions are unknown. *Management Science*, 33(4):542—-544, 1987.

Yin-Wong Cheung and Daniel Friedman. Individual learning in normal form games: Some laboratory results. *Games and Economic Behavior*, 19(1):46 – 76, 1997.

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.

Urs Fischbacher. z-Tree: Zurich toolbox for readymade economic experiments. *Experimental Economics*, 10(2):171–178, 2007.

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.

Ben Greiner. Subject pool recruitment procedures: organizing experiments with orsee. *Journal of the Economic Science Association*, 1(1):114–125, 2015.

Tanjim Hossain and Ryo Okui. The binarized scoring rule. *Review of Economic Studies*, 80(3):984–1001, 2013.

Kyle Hyndman, Antoine Terracol, and Jonathan Vaksmann. Strategic interactions and belief formation: an experiment. *Applied Economics Letters*, 17(17):1681–1685, 2010.

Kyle Hyndman, Erkut Y. Ozbay, Andrew Schotter, and Wolf Zefev Ehrblatt. Convergence: An experimental study of teaching and learning in repeated games. *Journal of the European Economic Association*, 10(3):573–604, 2012.

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.

Ehud Kalai and Ehud Lehrer. Rational learning leads to nash equilibrium. *Econometrica*, 61(5):1019–1045, 1993. ISSN 00129682, 14680262.

J. Keith Murnighan and Alvin E. Roth. Expecting continued play in prisoner's dilemma games. *Journal of Conflict Resolution*, 27(2):279–300, 1983.

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.

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. Mixed strategies in the indefinitely repeated prisoner's dilemma. Working Paper, 2019.

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.

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.