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

BELIEFS IN REPEATED GAMES

Masaki Aoyagi
Guillaume R. Fréchette
Sevgi Yuksel

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	Guillaume R. Fréchette	Sevgi Yuksel
Osaka University	NYU	UCSB

February 12, 2021

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 accurate despite some 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 leads to different beliefs, and the same belief leads to 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 overestimate the likelihood that others use the same strategy as them, while underestimating the likelihood that others use less cooperative strategies.

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, why does a variety of strategies with different levels of cooperation coexist in both finitely repeated and indefinitely repeated settings? Do these out-of-equilibrium phenomena (at least in light of simple standard models) stem from preferences, information, incorrect beliefs, or bounded rationality? In this paper, we bring to light one key force to help us better understand behavior in such an environment.

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. In repeated games, however, strategies are 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. A player of repeated games will have difficulty forming a belief that correctly predicts other players' strategies. This is potentially made even more difficult by multiplicity of equilibria. When repeated, the PD can generate diverse patterns of dynamic behavior rationalized by different beliefs, and thus provides an extremely informative framework for the joint study of beliefs and strategies. In this sense, making beliefs observable can establish facts that speak to how people reason in repeated games. For instance, we may find evidence pointing to non-standard preferences if strategies considered only from the perspective of payoffs do not best respond to beliefs, or to a failure of learning if the strategies do best respond to beliefs but beliefs are incorrect. Such evidence could explain the presence and persistence of cooperation in the finitely repeated PD and the variety of strategies observed in the indefinitely repeated PD.

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. In contrast to the theoretical predictions, experiments have shown certain game parameters can generate early cooperation in *both* environments. Eliciting beliefs will allow us to explore whether cooperation in these two canonical environments is sustained by similar forces, despite the theoretically distinct nature of the two games.

There is a rapidly growing literature on experiments with belief elicitation, but a large majority of these papers 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. The primary focus of these papers is the consistency of beliefs and actions, on which there are mixed results. We review this literature in detail in Appendix A. However, this literature is very different from the current paper. In particular, to the extent that such 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.[1] In contrast, our experiment elicits beliefs in repeated games where dynamic incentives are clearly important. As such, the only closely related paper is Gill & Rosokha (2020), who elicit beliefs from subjects playing multiple indefinitely repeated PD games. However, the paper differs from ours both in its focus and design. In Gill & Rosokha (2020), subjects directly choose strategies (from a list of 10) and in the first and last supergames also report beliefs over those strategies. Their design allows them to study beliefs from the onset, and as such their focus is more directly on the evolution of beliefs in indefinitely repeated games and their connection to a level-k model of rationality. They also link personality traits to strategies and beliefs. They find that strategy choices are broadly consistent with beliefs, which is also one of the findings in our experiment despite the important differences in our designs. They show this for a variety of indefinitely repeated games, we document this in both finitely and indefinitely repeated 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. Most importantly, we only elicit (first-order) beliefs about the other players' stage actions and not, for example, beliefs conditional on some action realization, beliefs over beliefs, or beliefs over supergame strategies. We also 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.

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

[1] 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.

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

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). 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 the Indefinite games, and provides insight into whether their strategic reasoning is similar or different across these two games.

We find beliefs covary with actions (on average and along many histories). However, we also document some systematic deviations that can be tied to the slow unravelling of cooperation in the Finite game. In addition, our results show beliefs are *not* simply the (weighted) empirical average of past observations and are *forward looking*. For example, subjects correctly anticipate a decline in the likelihood that their opponent will cooperate in later rounds of the Finite game.

The finding that beliefs are not simply the summary of past history implies the subjects are aware (or at least behave as such) of the possibility that their opponent is playing a supergame strategy that reacts to history in non-trivial ways. This suggests we need to look beyond beliefs over stage actions and consider *supergame beliefs*, defined here as beliefs over supergame strategies. In fact, the experimental literature hints at the possibility that supergame beliefs are a key driver of behavior in the repeated PD. For example, the data from previous repeated game experiments show evolution of behavior can be well described by learning and evolutionary models over supergames (Dal Bó & Fréchette 2018, Embrey et al. 2018, Proto et al. 2020).

Given that beliefs are elicited only on the realized path of play, supergame beliefs are not directly observable and need to be estimated. We propose a novel method to recover such beliefs: we first type subjects according to the supergame strategy they are estimated to be playing, and then estimate the supergame belief of each individual type separately.[3]

The results show that subjects who play different supergame strategies have different supergame beliefs. In fact, for many types, their (supergame) strategy is *subjectively rational* in the sense that it best responds to their supergame beliefs (and most come close to best responding).[4] These observations suggest heterogeneity in behavior can be explained, to a large extent, by heterogeneity in beliefs. More generally, in both the Finite and the

associated with the discount factor of the infinitely repeated game.

[3]This formulation is close to that explored in Kalai & 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.

[4]Note supergame beliefs and supergame strategies are estimated independently: the former uses belief reports and the latter uses action choices. Thus, the estimation does not necessarily imply such a relation.

Indefinite games, subjects' supergame beliefs correctly anticipate the type of strategies played in each environment, but are not necessarily well calibrated to the actual frequency of strategies in the population. Furthermore, supergame beliefs reveal a general tendency for subjects to overestimate the popularity of their own strategy and to underestimate the likelihood that others adopt less cooperative strategies than they do.

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.*[5] One interpretation of $\tilde{p}_i$ is that it represents i's prior belief over the

[5]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.

proportion of different strategies played by other subjects in that session.[6]

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

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

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

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

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

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

Some of the key supergame strategies in our analysis are as follows. AC and AD are the strategies that choose C and D, respectively, for every history. σ_i is Grim if $\sigma_i^t(h^{t-1}) = C$ if and only if $h^{t-1} = ((C, C), \ldots, (C, C))$. σ_i is TFT (resp. STFT) if $\sigma_i^1 = C$ (resp. $\sigma_i^1 = D$) and $\sigma_i^t(h^{t-1}) = a_j^{t-1}$ 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.[7]

[6] 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.

[7] All strategies considered in our analysis are listed and defined in Table 11 of Appendix B.

3 Design

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

1. *Selecting parameters such that initial cooperation rates are high in both the Finite and the Indefinite games (and at similar rates).* We can reasonably expect people to have different beliefs in the Finite and the Indefinite games, especially when these games generate very different behavior (cooperation rates). Although documenting this can have value, focusing on the more puzzling case where these two types of games generate very similar behavior (specifically in terms of initial cooperation rates) is potentially more interesting. Given the theoretical contrast between Finite and Indefinite games, beliefs can be particularly informative in providing insights on whether cooperation is driven by similar considerations across these two games. To achieve this, we based our parameter selections on previous experiments and meta-analysis (Embrey et al. 2018, Dal Bó & Fréchette 2018).

2. *Introducing belief elicitation while mitigating the impact the introduction might have on how subjects play.* One concern is that asking for beliefs from the onset of the experiment may alter how subjects approach the strategic interaction. To reduce this possibility, we separate the experiment into two parts. First, subjects are presented with "standard" repeated PD experimental instructions that do not mention beliefs. Second, after four supergames, the experiment is paused, and instructions explaining the belief-elicitation procedures are given. This two-part approach draws on Dal Bó & Fréchette (2019), who do this for strategy elicitation.[8] The results of the current experiment reproduce the qualitative features of previous experiments without belief elicitation (with similar parameters). Although beliefs at the start of the experiment are not elicited in the two-part approach, the potential benefits of not distorting behavior outweigh the downside of not observing beliefs in those early supergames.

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 & Fréchette 2018).[9] For instance, for the parameters we use in the Finite game, Embrey et al. (2018) find the average round of the last cooperation moves one round earlier for every 10 supergames. 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.

[8]They find that choices in their experiments with strategy elicitation are similar to those from experiments without strategy elicitation.

[9]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.

Table 1: Stage Game

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

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.[10] We use *supergame* to refer to each repeated game played between two matched players, and *round* to refer to each play of the stage game. In the Finite game, each supergame ends after eight rounds, $T = 8$.[11] In the Indefinite game, after each round, there is a $\frac{7}{8}$ probability that the supergame will continue for an additional round.[12] 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.[13] In the Indefinite game, *observation rounds* refer to the rounds in which the subjects actually made action choices, and *game rounds* refer to those rounds that were part of the supergames. We denote by T the number of observation rounds in the Indefinite game so that $T = \max\{8, \text{“No. of game rounds”}\}$. For example, if an Indefinite game has five rounds, $T = 8$ because we observe the subject make eight choices even though only the first five mattered for payoffs, whereas if a supergame lasts 10 rounds, $T = 10$.

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. 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.[14] 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

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

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

[12] 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.

[13] This method was first introduced in Fréchette & Yuksel (2017) and has now been used in multiple papers on a variety of topics, for example, Vespa & Wilson (2019) in dynamic games, Agranov et al. (2016) in bargaining, and Weber et al. (2018) in a bond market.

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

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: Late	Total no. of Obs. Rounds
	1	20	12			8, 8,		96
	2	20	12			8, 8,		96
	3	20	13			8, 8, 8,		104
Finite	4	20	11	8, 8, 8, 8	*8, 8, 8*	8,	**8, 8, 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
	1	20	10	9, 7, 13, 7	*1, 2, 23,*		**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
Indefinite	4	16	8	9, 7, 10, 13	*32,*		**7, 7, 6**	96
	5	14	12	7, 22, 7, 3	*2, 5, 8,*	4, 14,	**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).

choice of the matched subject.[15] The belief question is presented on a separate screen after subjects have made their action decision for that round and before feedback is provided. 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 & 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 & Filippin 2013). Instructions for this task were distributed after the completion of the last supergame.[16]

We conducted eight sessions per treatment and 16 sessions in total.[17] Table 2 summarizes basic information about each session. The supergames for the part with belief

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

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

[17] 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.

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.[18] 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.[19]

4 Results

The analysis of our data is separated into three sections. Section 4.1 provides an 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. Finally, 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

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

[18] 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.

[19] 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 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)).

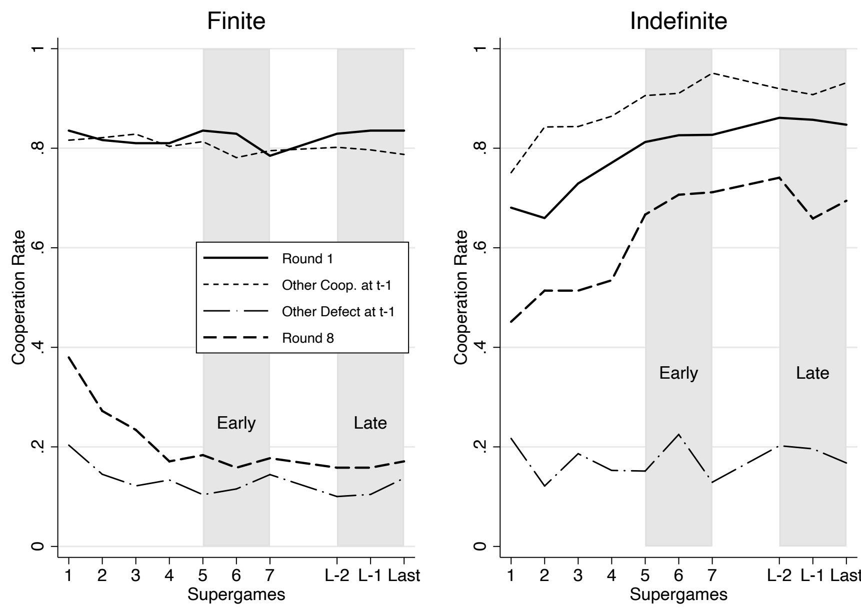

Figure 1: Cooperation Rate over Supergames

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.[20]

Hence, consistent with prior experiments with similar 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, and we find no indication of important changes in the subjects' behavior caused by the belief-elicitation task in these environments.[21]

Result 1 We reproduce qualitative data patterns observed in previous experiments on

[20]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 Appendix B (Figure 13).

[21]Table 6 in the Appendix B also shows no significant changes in round-one choices for supergames where beliefs are elicited (compared with those where they are not).

Finite and Indefinite PD games. In particular, our results confirm cooperation is history dependent in both games. Furthermore, cooperation evolves differently in both games: it collapses at the end only in the Finite game.

4.2 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.

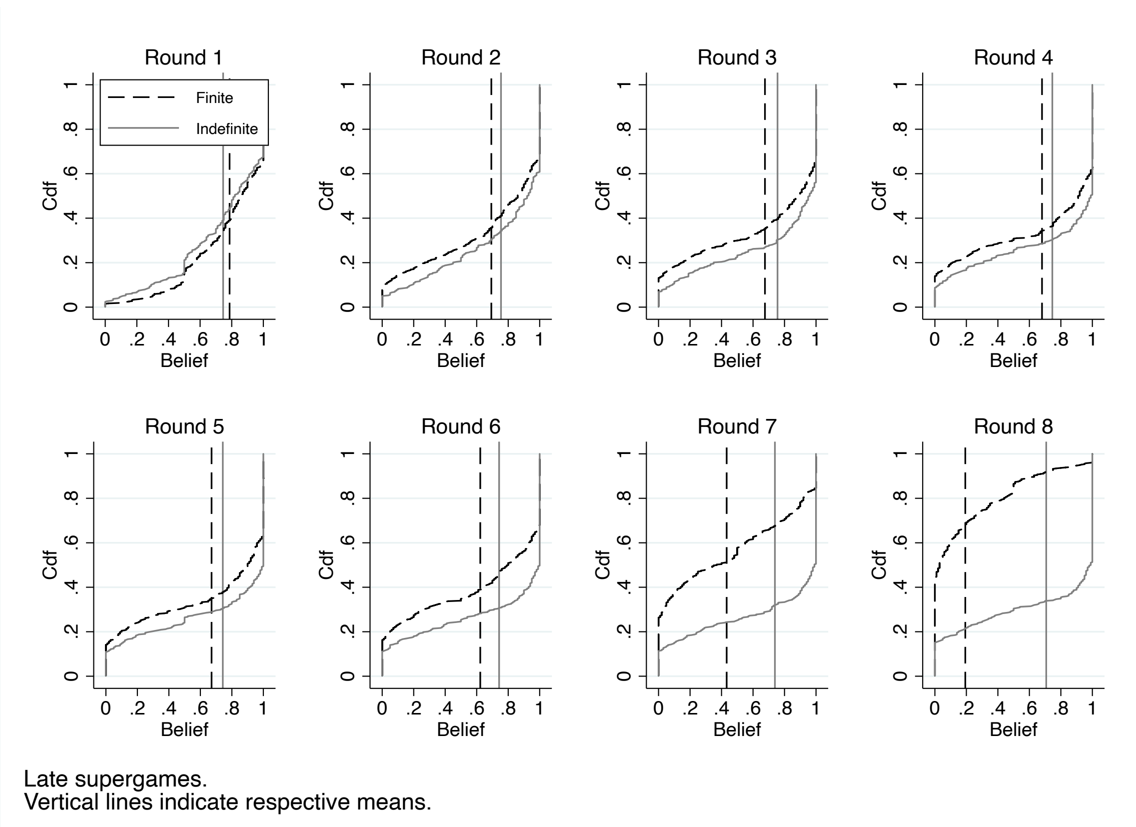

Figure 2: Distribution of Beliefs by Round

Figure 2 displays the cumulative distributions of (unconditional) beliefs in rounds $t = 1, \ldots, 8$ of late supergames in the Finite and Indefinite games.[22] As the figure clearly

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

shows, beliefs evolve very differently over rounds across these two types of games.[23] Beliefs become comparatively more pessimistic in the Finite game as the supergame unfolds. The difference in the average belief is statistically significant in rounds six, seven, and eight.[24]

A few more observations are worth making. First, beliefs are varied and do not concentrate on a few values. Second, subjects do report beliefs of 0 and 1, not just interior values. In fact, in round eight, more than 40% of subjects place probability one on defection by the other player in the Finite game, whereas more than 40% of subjects place probability one on cooperation by the other player in the Indefinite game.

Result 2 Beliefs are different in Finite and Indefinite games. The main difference is that beliefs about cooperation collapse toward the end in the Finite game.

4.2.1 Actions and Beliefs

Putting beliefs and actions together reveals beliefs–on average–track cooperation rates closely. Figure 3 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 Indefinite. 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 magnitude).[25]

However, when we look at each round separately, both in the Finite and the 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

[23]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.

[24]Respectively, $p < 0.05$, $p < 0.01$, and $p < 0.01$. 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.

[25]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.

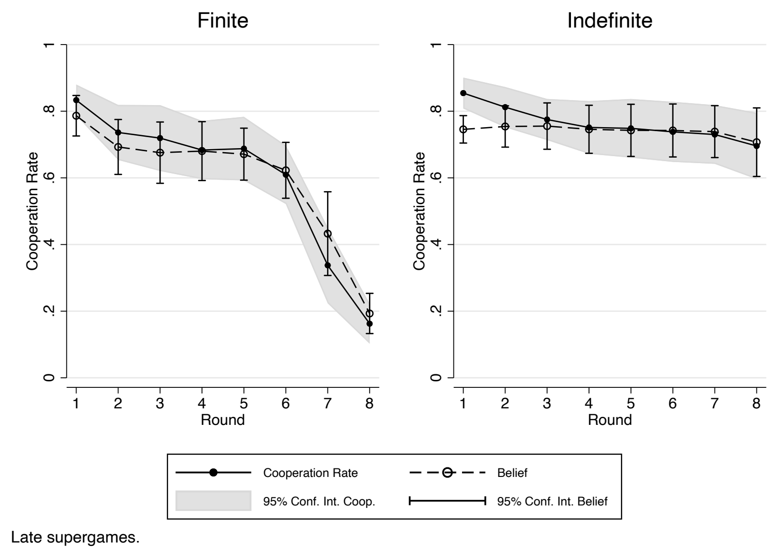

Figure 3: Choices and Beliefs by Round

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 mis-predictions. Figure 4 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 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 interest. One conjecture is that subjects are simply reporting back their observations about others' behavior from previous supergames. Alternatively, subjects may form beliefs relying on

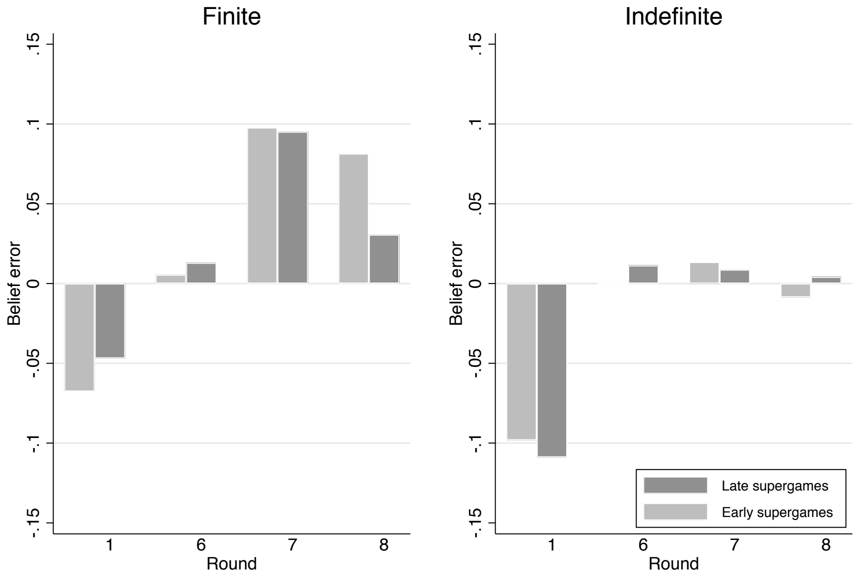

Belief error denotes average difference between beliefs and actions.

Figure 4: Belief Errors in Early vs. Late Supergames

introspection alone, or some combination of learning and introspection.[26] The data suggests that although experiences matter in shaping beliefs, they are not the sole determinant. Figure 16 in 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. Although each panel displays a peak close to 0, many are relatively flat and some are not centered at zero.

So far in Figures 3 and 4, 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 5 and 6 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.[27] Interestingly, note the downward adjustments following a unilateral choice of D by either player in round one are

[26] 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.

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

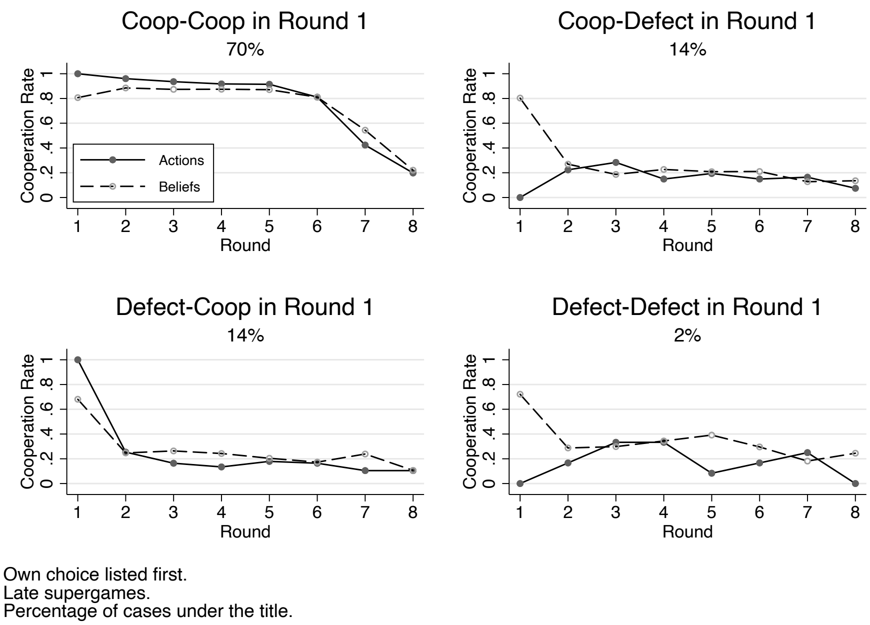

Figure 5: Conditional Round-Two Beliefs, Finite Games

of the correct magnitude even though the required adjustment is quite large. Comparing the two figures, we see action frequencies and beliefs evolve in a similar fashion in all panels except for the top-left panel, which shows clear differences across the two treatments. In the Finite game, most of the initially cooperative interactions eventually break down, and 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.

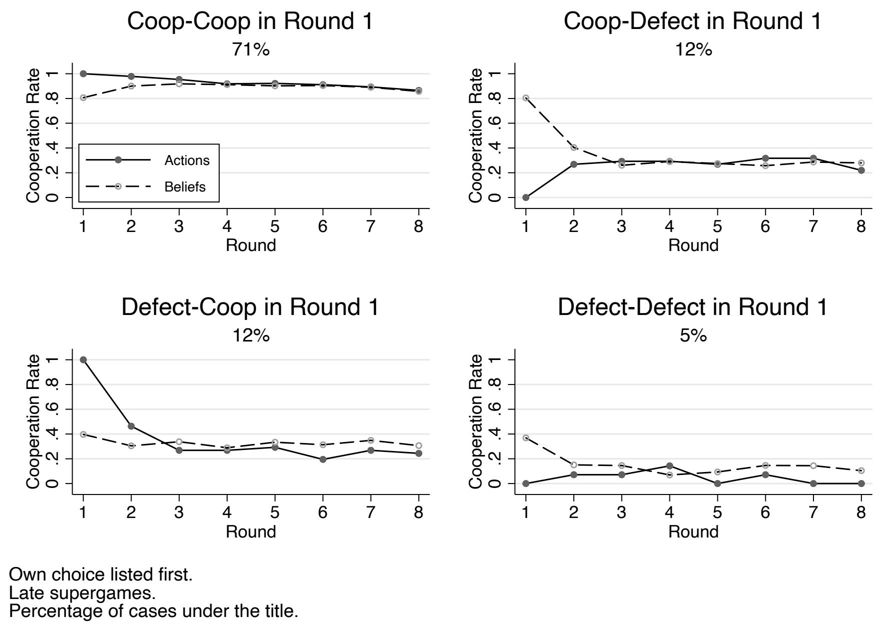

Figure 6: Conditional Round-Two Beliefs, Indefinite Games

Table 8 in 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) even in early supergames. 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 8 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).

As Figures 5 and 6 above show, supergames starting with joint cooperation are the most common. How do beliefs evolve on a mutual cooperation path? Figure 7 shows the average cooperation rates $\bar{x}^t$ and average beliefs $\bar{\mu}^t$ along the history $h^{t-1} = ((C, C), \ldots, (C, C))$.[28]

[28]Note $t = 2$ corresponds to cases presented in Figures 5 and 6.

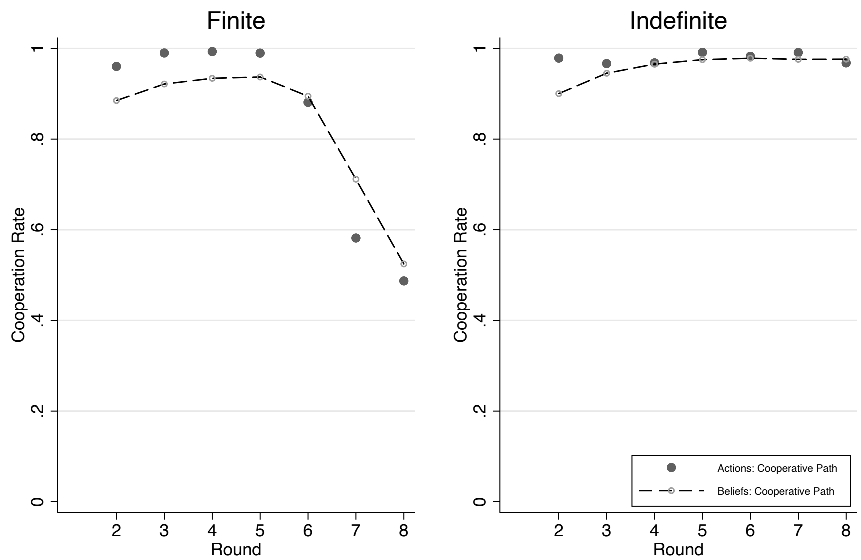

Cooperative path: joint cooperation up to that round.
Late supergames.

Figure 7: Cooperative Path (First Eight Rounds)

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.[29] 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.[30] 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 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

[29]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.

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

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.[31] Clearly, beliefs in the Finite game change on a cooperative path with the length of the interaction, and hence are non-stationary.

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

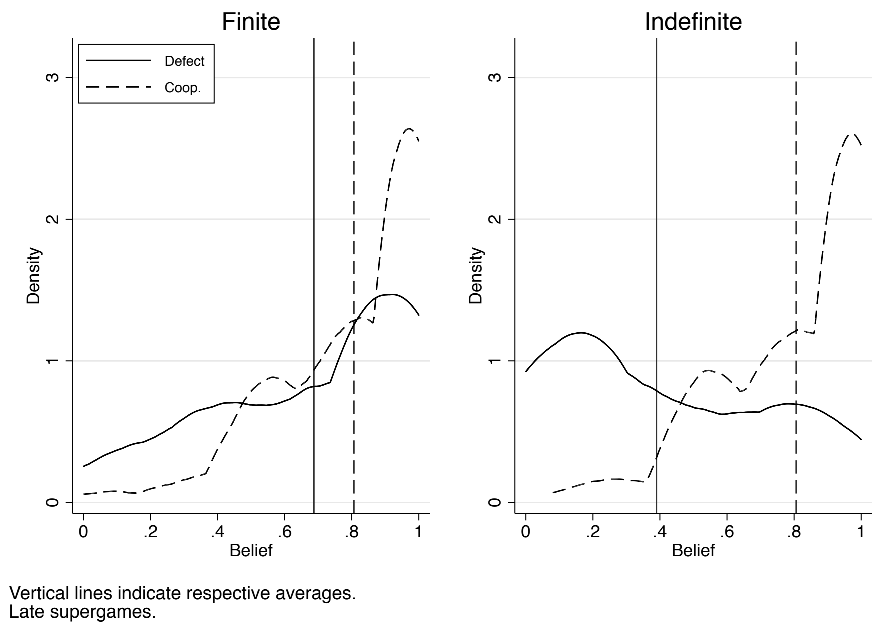

Figure 8: Beliefs of Defectors vs. Cooperators in Round One

We now turn to the question of whether different actions are supported by different beliefs. Figure 8 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

[31]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 $\frac{1}{n}\sum_i|x_{m(i)}^6 - x_{m(i)}^7|$.

and defectors are statistically different in the Indefinite game ($p < 0.01$) but not in the Finite game.[32] 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. 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 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.[33]

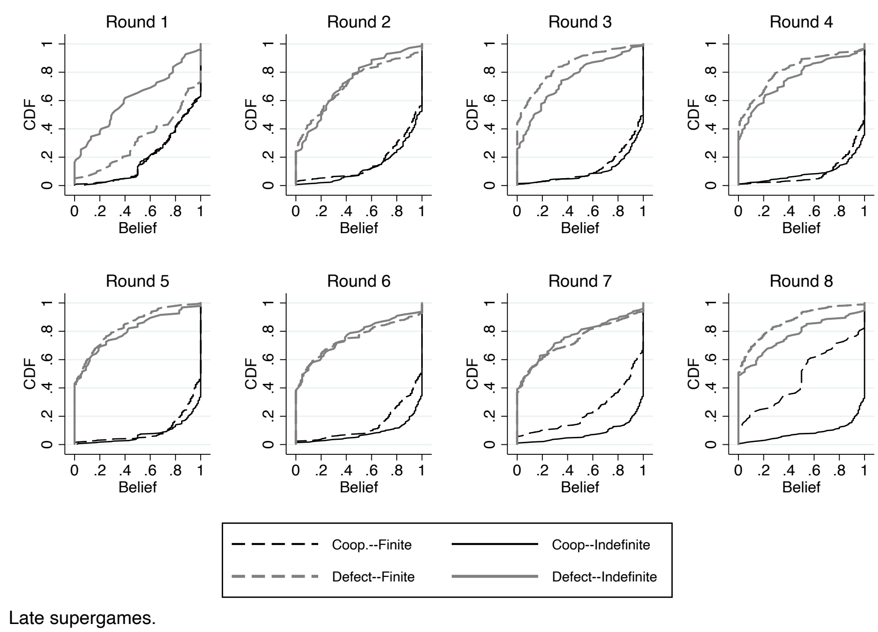

Figure 9: Beliefs by Action and Treatment: Rounds One through Eight

Figure 9 plots the CDF of beliefs by action and treatment for each round. It clearly

[32]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.

[33]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.

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 9 in Appendix B shows that in all rounds, the marginal impacts of beliefs on the likelihood of cooperation are 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.[34] More specifically, subjects are more likely to defect later in the Finite game, even if their beliefs are the same.

Perhaps more surprisingly, 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 4 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.[35] The estimation method we develop has three 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. Estimate strategies at the population level.
2. Use these estimates and each subject's choices to classify them into types.
3. Estimate beliefs over supergame strategies separately for each type.

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

[35] We do not make the claim that subjects reason in terms of strategies *per se*, but that we can represent their behavior as such.

Details are provided in sections 4.3.1, 4.3.2, and 4.3.3 below. 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 considers 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. How do we specify the relevant set of strategies for our data set? By now, a significant body of literature documents which strategies are used in repeated PD games. 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. This forces us to 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.[36]

[36] To explore whether subjects of the same type have similar beliefs, 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

4.3.1 Population-Level Estimates of Strategies

We first use the Strategy Frequency Estimation Method (SFEM) introduced in Dal Bó & Fréchette (2011) to estimate the distribution of strategies used.[37] 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. Then, we use results on this set to narrow our focus to 10 strategies. The larger 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 perfect monitoring.[38] 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.[39] Results on this consideration set are reported in Appendix B. However, because our primary goal is to estimate beliefs over strategies, focusing on such a large set is more costly than it is typically 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.[40]

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

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.

[37] See Dal Bó & Fréchette (2018) for results across a variety of experiments applying the SFEM to indefinite PDs with perfect monitoring. Results from direct elicitation of strategies (Dal Bó & Fréchette (2019) and Romero & Rosokha (2019)) are in line with SFEM results and indicate that for indefinite PDs with perfect monitoring, the majority of subjects use strategies with a TFT, Grim, or AD structure.

[38] Our aim was to be inclusive in the first step of the selection process. A strategy is included if it has been shown to have predictive power in *any* setting among a variety of environments covering a large parameter set.

[39] 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.

[40] 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: 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.

account for a majority of the strategies in the Indefinite game (Dal Bó & Fréchette 2018).[41]

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. 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 5 We reproduce results about strategy choices observed in previous finitely and indefinitely repeated PD games.[42] In particular, our results confirm strategic heterogeneity exists within and across treatments. In the Finite game, subjects mostly use threshold

[41]The Appendix also reports SFEM results for early supergames (the changes are presented in Figure 19). Consistent with Embrey et al. (2018) those results show that threshold strategies increase with experience in the Finite game.

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

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. Recall the SFEM yields the probability distribution over supergame 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.[43]

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.[44] 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 $\frac{p_{\hat{AD}}\beta^4(1-\beta)^{20}}{p_{\hat{AD}}\beta^4(1-\beta)^{20}+p_{\hat{AC}}\beta^{20}(1-\beta)^4}$, 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.[45] Clearly,

[43] A unique strategy exists within the consideration set for each subject in our data set that achieves the highest posterior (given the SFEM results).

[44] One could use a different prior. We have explored using a uniform prior in simulations, and the results are far worse than with the SFEM estimates.

[45] 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

the smaller the fraction of subjects of a given type, the less reliable their belief estimates will be.

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.[46] 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 $\tilde{Z}^{\beta}$ 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 Online 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

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 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.

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

three of the 84 estimates 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 15.

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 15.

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

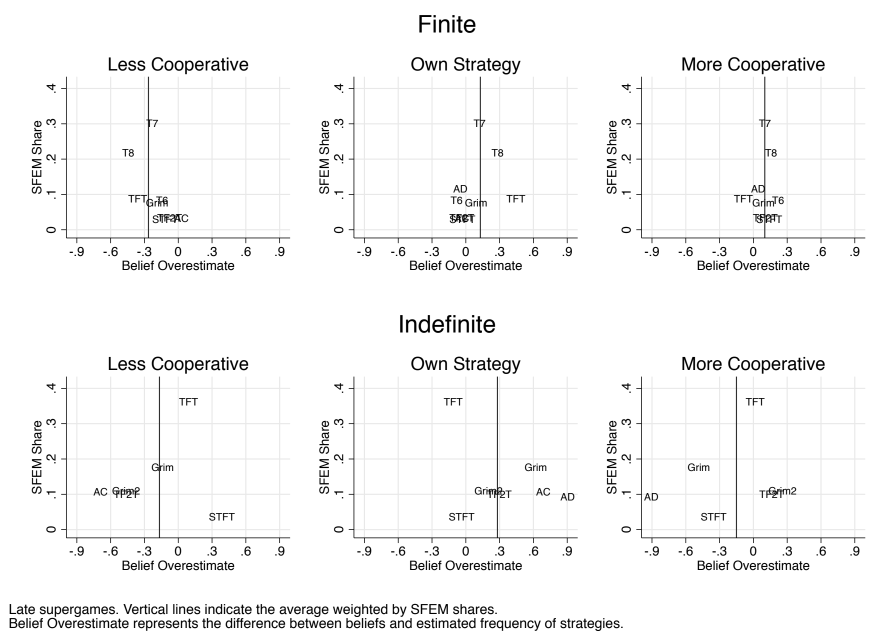

Figure 10: Belief vs. Estimated Frequencies

Tables 4 and 5 report detailed information on how beliefs differ by type in each game.

However, the richness of these data make identifying general patterns on how beliefs over strategies are connected to subjects' strategy choice difficult. To simplify our analysis, we order strategies in terms of *cooperativeness* and use this order to study the degree to which subjects believe others are *less*, *as*, or *more* cooperative than they are. 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.[47] This generates the following order of cooperativeness (from least to most): AD, STFT, T6, T7, 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. Figure 10 makes use of this categorization to study the accuracy of beliefs about others being *less*, *as*, or *more* cooperative. Each of these cases is presented in a separate panel. In each graph, the x-coordinate for each type represents the difference between beliefs and the estimated frequency of strategies. Consider, for example, T8 in the lop left panel. First, we compute T8's subjective belief about the likelihood that others are using a less cooperative strategy than T8. This involves summing over T8's beliefs about others using AD, STFT, T6 or T7, which are the only strategies less cooperative than T8. From this we subtract the objective likelihood (obtained from the SFEM) of others using AD, STFT, T6 or T7. For T8, we see that this difference is negative. This means that T8 types underestimate the likelihood that others are less cooperative than they are. On these graphs, the y-coordinate for each type represents the estimated frequency of the type, allowing us to evaluate the prevalence of such deviations. The figure suggests a tendency shared in both games: (1) Subjects underestimate (relative to actual frequencies) the likelihood that others use less cooperative strategies than their own. (2) Subjects overestimate the likelihood that others use the same strategy as their own. The most pronounced between these distortions is the first for the Finite game and the second for the Indefinite game. With respect to beliefs on others being more cooperative, we observe opposite patterns in the Finite and the Indefinite games. In the Finite game, subjects overestimate the likelihood of such an event; in the Indefinite game, subjects underestimate it. However, in both cases, this is the smallest of the three distortions. Note also that the most popular type in the Finite game (T7—as well as the second most popular T8) displays the tendency observed, on average, whereas for the Indefinite game, TFT (the most popular type) is well calibrated.

The accuracy of beliefs over strategies can be studied more directly without relying on the cooperativeness order. In Appendix B, we compute, for each type, the Euclidean distance between beliefs and the estimated frequency of strategies. To study whether beliefs

[47]Analytical derivation of the cooperativeness order for an infinitely repeated PD with a general specification of the payoffs and discount factor is presented in Appendix B. It also describes 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.

become more accurate with experience, we also look at how this distance changes from early to late supergames.[48] 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.[49]

Result 7 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.

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 their strategy choice is to best responding to their beliefs. 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 call a type subjectively rational (given preferences induced in the experiment) if her strategy choice is a best response to her supergame beliefs among those strategies Z in the consideration set.[50]

The results, presented in Figures 11 and 12, suggest most subjects' strategy choices are either exact or approximate best responses given their supergame beliefs.[51] 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.[52]

[48] Additionally, in the same Appendix, we document in detail how the distribution of strategies, types, and beliefs for each type change from early to late supergames.

[49] 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.

[50] 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 $\tilde{\beta}$.

[51] 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.

[52] 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.

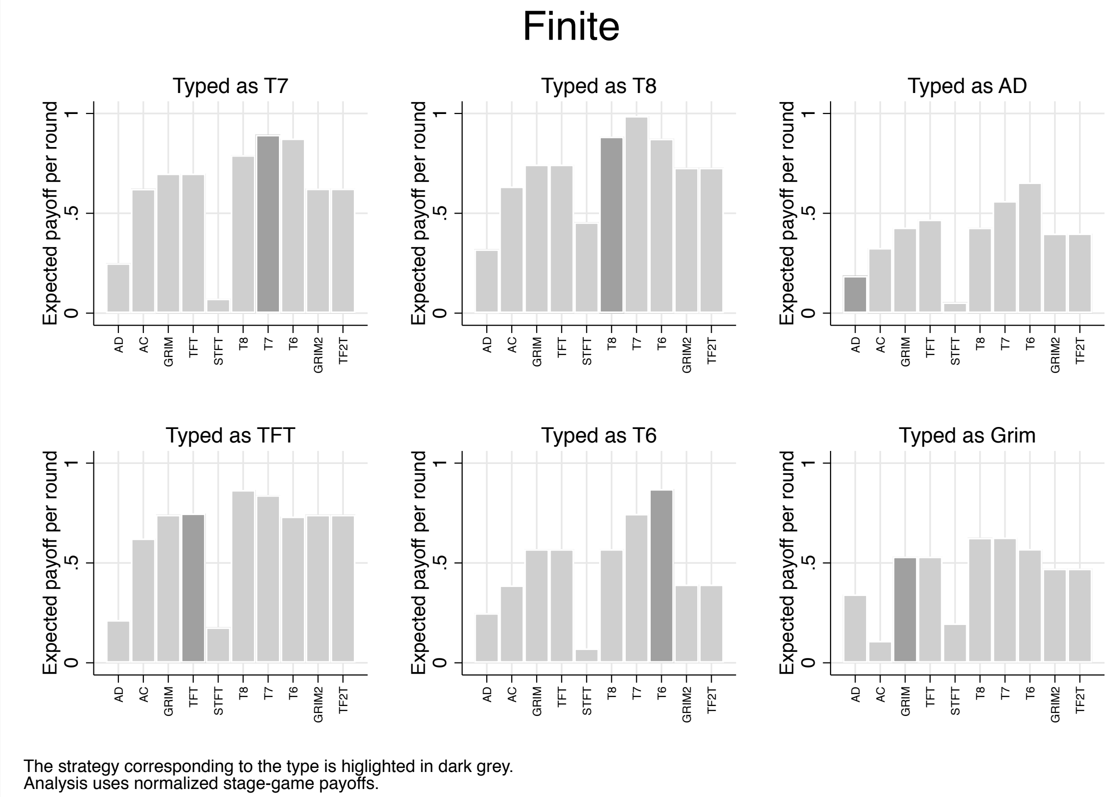

Figure 11: Best Response for Top 6 Types in the Finite Game

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.[53] 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.[54] Hence, overall, the majority of subjects appear subjectively rational or close to subjectively rational.

Result 8 Most types are close to best responding to their beliefs: they are subjectively rational.

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

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

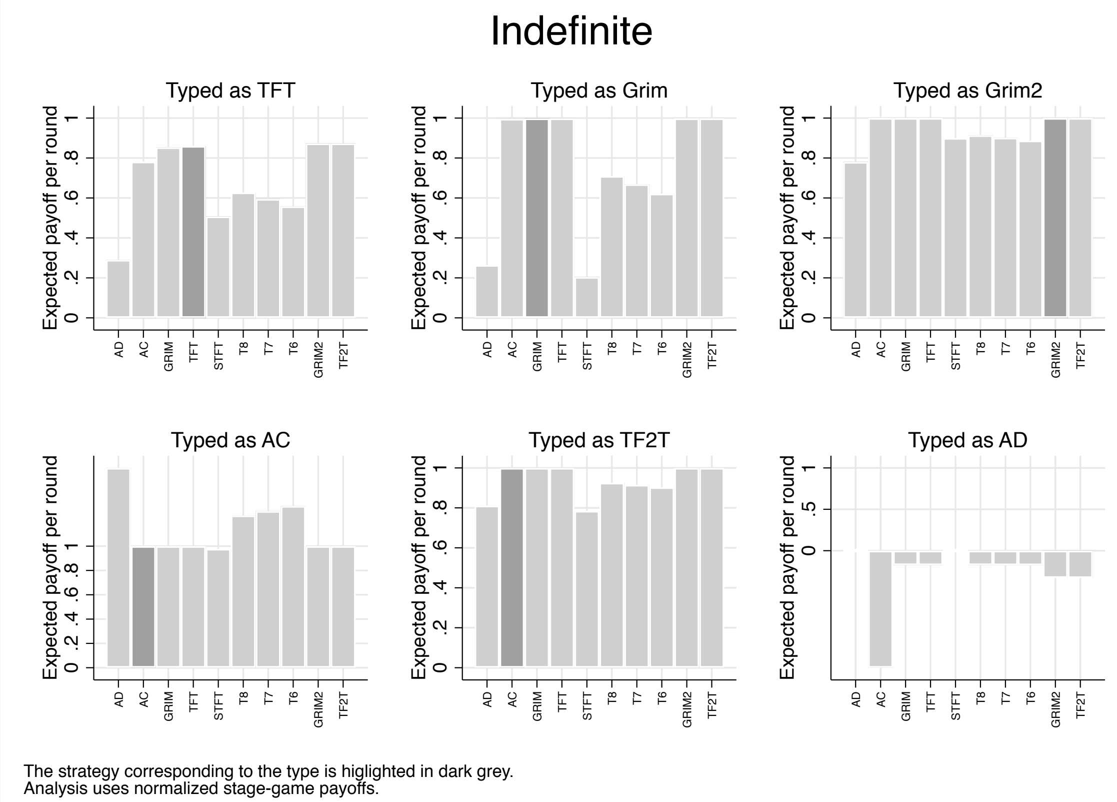

Figure 12: Best Response for Top 6 Types in the Indefinite Game

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, we see the most frequent TFT type achieves 94% of the best-response payoff from Grim2. However, some strategy-types are further away from best responding to the population. For example, the AD type in the Finite game only achieves 64% of the best-response payoff.

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.

Our first key finding is that 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. To the extent that average beliefs do not correspond to the empirical action frequencies, the key deviations are over-optimism in late rounds of the Finite game and over-pessimism in early rounds of the Indefinite game. 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. On the individual level, although beliefs are heterogeneous with a varying degree of accuracy, a vast majority of subjects correctly anticipate the likelihood of their opponent cooperating.

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 again correctly anticipated along such histories. In general, beliefs are different across the two games even when they lead to the same action.

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, we find subjects in the Finite and the Indefinite games correctly anticipate the different class of strategies used in these games: threshold strategies in the former and conditionally cooperative strategies in the later. Second, we study the connection between the choice of supergame strategies and beliefs over them. We observe that, in both the Finite and the Indefinite games, subjects playing different strategies have strikingly heterogeneous beliefs over the strategy choice of the other player. In both games, the results also suggest a tendency by subjects to overestimate the likelihood that others use the same strategy as they do, while underestimating the likelihood that others use less cooperative strategies. Third, when we study the link between strategies and beliefs, we find most types are close to being subjectively rational; that is, we observe the expected payoff associated with their strategy choice to be close to their best-response payoff given their beliefs. This, in particular, suggests the observed heterogeneity in strategy choice is related to the heterogeneity in beliefs.

These results also provide insights into the forces that underlie some of the key behavioral patterns observed in these games. Our results illustrate 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 fairly accurate, for 80% of subjects, best responding to their beliefs involves cooperating more than it is optimal (given the observed strategy distribution in the population). 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.

References

Agranov, M., Cotton, C. & Tergiman, C. (2016), 'Persistence of power: Repeated multilateral bargaining', *Journal of Public Economics* . forthcoming.

Allen, F. (1987), 'Discovering personal probabilities when utility functions are unknown', *Management Science* **33**(4), 542–544.

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

Chaudhuri, A., Paichayontvijit, T. & Smith, A. (2017), 'Belief heterogeneity and contributions decay among conditional cooperators in public goods games', *Journal of Economic Psychology* **58**, 15 – 30.

Cheung, Y.-W. & Friedman, D. (1997), 'Individual learning in normal form games: Some laboratory results', *Games and Economic Behavior* **19**(1), 46 – 76.

Costa-Gomes, M. A., Huck, S. & Weizsäcker, G. (2014), 'Beliefs and actions in the trust game: Creating instrumental variables to estimate the causal effect', *Games and Economic Behavior* **88**, 298 – 309.

Costa-Gomes, M. A. & Weizsäcker, G. (2008), 'Stated Beliefs and Play in Normal-Form Games', *The Review of Economic Studies* **75**(3), 729–762.

Crosetto, P. & Filippin, A. (2013), 'The "bomb" risk elicitation task', *Journal of Risk and Uncertainty* **47**(1), 31–65.

Dal Bó, P. (2005), 'Cooperation under the shadow of the future: experimental evidence from infinitely repeated games', *American Economics Review* **95**(5), 1591–1604.

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

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

Dal Bó, P. & Fréchette, G. R. (2019), 'Strategy choice in the infinitely repeated prisoners dilemma', *American Economic Review* . Forthcoming.

Danz, D. N., Fehr, D. & Kübler, D. (2012), 'Information and beliefs in a repeated normal-form game', *Experimental Economics* **15**(4), 622–640.

Danz, D., Vesterlund, L. & Wilson, A. J. (2020), Belief elicitation: Limiting truth telling with information on incentives, Technical report, National Bureau of Economic Research.

Embrey, M., Fréchette, G. R. & Yuksel, S. (2018), 'Cooperation in the finitely repeated prisoner's dilemma', *The Quarterly Journal of Economics* **133**(1), 509–551.

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

Fischbacher, U. & Gächter, S. (2010), 'Social preferences, beliefs, and the dynamics of free riding in public goods experiments', *American Economic Review* **100**(1), 541–56.

Flood, M. M. (1952), 'Some experimental games', *Research Memorandum RM-789, The Rand Corporation, Santa Monica* .

Fréchette, G. R. (2012), 'Session-effects in the laboratory', *Experimental Economics* **15**(3), 485–498.

Fréchette, G. R. & Yuksel, S. (2017), 'Infinitely repeated games in the laboratory: Four perspectives on discounting and random termination', *Experimental Economics* **20**(2), 279–308.

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

Gächter, S. & Renner, E. (2010), 'The effects of (incentivized) belief elicitation in public goods experiments', *Experimental Economics* **13**(3), 364–377.

Gill, D. & Rosokha, Y. (2020), Beliefs, learning, and personality in the indefinitely repeated prisoner's dilemma. Working Paper.

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

Hossain, T. & Okui, R. (2013), 'The binarized scoring rule', *Review of Economic Studies* **80**(3), 984–1001.

Hyndman, K., Özbay, E. Y., Schotter, A. & Ehrblatt, W. (2012*a*), 'Belief formation: an experiment with outside observers', *Experimental Economics* **15**(1), 176–203.

Hyndman, K., Ozbay, E. Y., Schotter, A. & Ehrblatt, W. Z. (2012*b*), 'Convergence: An experimental study of teaching and learning in repeated games', *Journal of the European Economic Association* **10**(3), 573–604.

Hyndman, K., Terracol, A. & Vaksmann, J. (2010), 'Strategic interactions and belief formation: an experiment', *Applied Economics Letters* **17**(17), 1681–1685.

Kalai, E. & Lehrer, E. (1993), 'Rational learning leads to nash equilibrium', *Econometrica* **61**(5), 1019–1045.

Kocher, M. G., Martinsson, P., Matzat, D. & Wollbrant, C. (2015), 'The role of beliefs, trust, and risk in contributions to a public good', *Journal of Economic Psychology* **51**, 236 – 244.

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

Neugebauer, T., Perote, J., Schmidt, U. & Loos, M. (2009), 'Selfish-biased conditional cooperation: On the decline of contributions in repeated public goods experiments', *Journal of Economic Psychology* **30**(1), 52 – 60.

Nyarko, Y. & Schotter, A. (2002), 'An experimental study of belief learning using elicited beliefs', *Econometrica* **70**(3), 971–1005.

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

Proto, E., Rustichini, A. & Sofianos, A. (2020), Intelligence, errors and strategic choices in the repeated prisoners' dilemma. Working Paper.

Rey-Biel, P. (2009), 'Equilibrium play and best response to (stated) beliefs in normal form games', *Games and Economic Behavior* **65**(2), 572 – 585.

Romero, J. & Rosokha, Y. (2019), Mixed strategies in the indefinitely repeated prisoner's dilemma. Working Paper.

Smith, A. (2013), 'Estimating the causal effect of beliefs on contributions in repeated public good games', *Experimental Economics* **16**(3), 414–425.

Smith, A. (2015), 'Modeling the dynamics of contributions and beliefs in repeated public good games', *Economics Bulletin* **35**, 1501–1509.

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

Weber, M., Duffy, J. & Schram, A. (2018), 'An experimental study of bond market pricing', *The Journal of Finance* **73**(4), 1857–1892.

Wilson, A. J. & Vespa, E. (2018), Paired-uniform scoring: Implementing a binarized scoring rule with non-mathematical language. Working Paper.

Online Appendix for

Beliefs in Repeated Games

Masaki Aoyagi Guillaume R. Fréchette Sevgi Yuksel

CONTENTS:

A Related Literature

Nyarko & Schotter (2002) are among the first to study elicited beliefs in repeated games. Studying a finite repetition of a 2×2 game with a unique mixed Nash equilibrium (NE) played in fixed and random pairing, Nyarko & Schotter (2002) find the subjects' beliefs over actions are not empirical in the sense that they cannot be approximated by the weighted average of the opponent's past actions.[55] Hyndman et al. (2010) study beliefs when a stage game with a unique mixed NE is repeated 20 times, and find subjects' beliefs about the other's action in the present round do take into account the effect of their own action choice in the preceding rounds, and hence cannot be expressed by the weighted average of the other player's actions in the past. Hyndman et al. (2012*b*) advance this observation in an experiment in which subjects play a finite repetition of 3×3 and 4×4 normal form games with and without dominance-solvable NE. Hyndman et al. (2012*b*) note some players attempt to influence the beliefs of other players through their own actions, and thus help the process converge to an NE.[56]

The experimental literature on beliefs examines the question of whether actions are best responses to beliefs with no definite answers. Nyarko & Schotter (2002) find the actions in each round mostly best respond to the stated beliefs, but also find fictitious-play beliefs better predict the opponents' action than the stated beliefs. Costa-Gomes & Weizsäcker (2008) use 14 3×3 games to investigate the relationship between subjects' elicited beliefs and their strategy choice. Regardless of whether belief elicitation precedes strategy choice, Costa-Gomes & Weizsäcker (2008) find the strategies are not best responses to the beliefs in a half of the games, and attribute this finding to the difference in the perception of the game in the two situations. Danz et al. (2012) use a dominance-solvable 3×3 game repeated 20 times to study beliefs under various combinations of feedback and matching conditions. Danz et al. (2012) find feedback of past actions helps advance the iterative elimination process both in terms of actions and beliefs. Using a series of 3×3 games each with a unique NE, Rey-Biel (2009) find more than two-thirds of subjects choose actions that best respond to their elicited beliefs.

The literature on voluntary-contribution games often finds conditional cooperation, which refers to the fact that subjects make higher contributions if they believe other members of their group make higher contributions. This relationship is observed, for example by Gächter & Renner (2010), Fischbacher & Gächter (2010) and Kocher et al. (2015).[57] Neugebauer et al. (2009) confirm this relationship in their experiment on a finitely repeated

[55] Nyarko & Schotter (2002) specifically consider a generalization of fictitious play called the γ-empirical average as proposed by Cheung & Friedman (1997).

[56] Hyndman et al. (2012*a*) have outside observers predict the actions of the subjects in Hyndman et al. (2012*b*), and find a large variance in their beliefs both in terms of accuracy and updating.

[57] Costa-Gomes et al. (2014) analyze the relationship in the trust game. Smith (2013, 2015) note the beliefs are endogenous, and hence that the effect on contribution, if interpreted as causal, is overestimated.

voluntary contribution game, and further observe that both contribution levels and beliefs about others' contribution levels decline toward the end. Chaudhuri et al. (2017) observe similar joint dynamics of contribution and beliefs, allowing for heterogeneity across subjects and classifying them into types according to their initial beliefs about others' contributions.

On cooperation and strategies in finitely and infinitely repeated PD, Dal Bó & Fréchette (2018) and Embrey et al. (2018) find some key patterns by re-analyzing data from a collection of laboratory experiments.[58] First, in finitely repeated PD, the fraction of threshold strategies increases with experience.[59] By the end, threshold strategies always account for the majority of the data, and use of the threshold strategies with lower thresholds increases with experience. This contributes to a (sometimes very) slow aggregate movement toward earlier defection.[60] In the finitely repeated PD, if the parameters are conducive to cooperation, round-one cooperation increases with experience, whereas last-round cooperation decreases with it.[61] Otherwise, cooperation remains low in all rounds. In indefinitely repeated PD, on the other hand, experience leads cooperation (in the first or last round) to almost any level, depending on how conducive the parameters are to cooperation. Experience also amplifies the magnitude of the effects of the parameters, although it does not change the direction of those effects. In most experiments with perfect monitoring, a few simple strategies account for more than 50% of the strategies used. They are "always defect" (AD), "always cooperate" (AC), "grim trigger" (Grim), "tit-for-tat" (TFT), and "suspicious-tit-for-tat" (STFT).[62] AD, Grim, and TFT are generally the most popular, and Grim becomes more popular with experience and appears to be a counterpart to the threshold strategies in finite games. The implementation error term in Grim also decreases with experience.[63] Experience also increases *responsiveness*, which is measured as the difference between the probability of cooperative action after cooperation by the other player and that after defection by the other player. This is documented in Aoyagi et al. (2019) and confirmed by Dal Bó & Fréchette (2018) in their analysis of the meta-data and new experiments: according to a simple regression, experience has a significant positive impact on responsiveness in 11 paper-treatments, whereas it is insignificant in 20 paper-treatments.[64]

[58] Experimental research on the subject goes as far back as Flood (1952).

[59] A threshold strategy (with threshold $k \geq 2$) starts with C and plays like grim-trigger before round k, but reverts to the unconditional play of D from round k on.

[60] Embrey et al. (2018) find that in the treatment most conducive to cooperation (replicated by the finite treatment of this study), the modal round at which cooperation breaks down moves earlier approximately by one round every 10 supergames.

[61] A longer horizon T, a higher discount factor δ, a lower temptation payoff $1 + g$, or a higher sucker payoff $-\ell$ all induce more cooperation.

[62] Grim cooperates until a defection is observed, at which point it defects forever; TFT starts by cooperating and thereafter matches what the other did in the previous round; STFT starts by defecting and thereafter matches what the other did in the previous round.

[63] See Dal Bó & Fréchette (2019), Tables 8 and A10.

[64] This analysis eliminates all data in within-subjects designs after a change in treatment and only preserves the initial treatment. Most of the insignificant cases have a small number of observations. One

treatment sees a negatively significant impact, perhaps because of relatively low round-one cooperation at 0.36.

B Additional Figures and Analysis

B.1 The Indefinite Game: Round Eight, Last Game Round, Last Observation Round

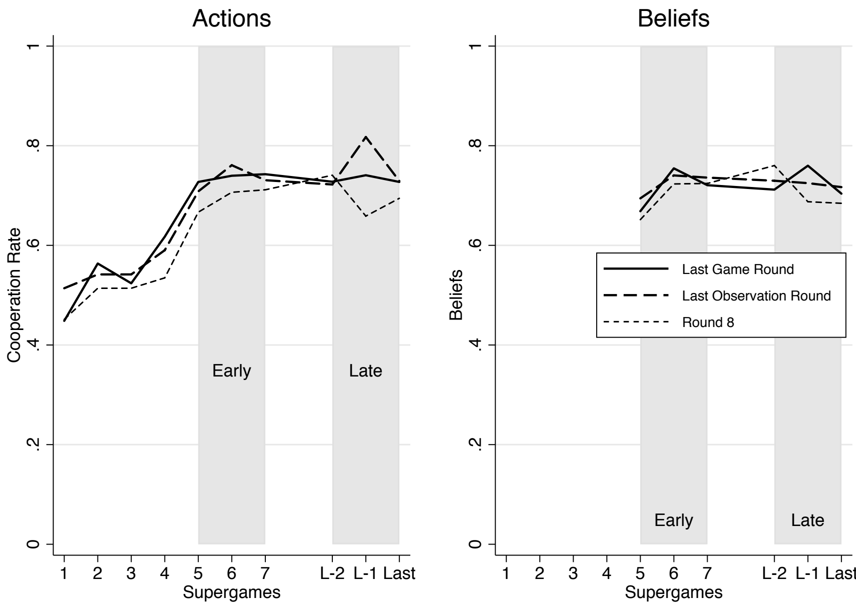

Figure 13: The Indefinite Game

B.2 Effect of Belief Elicitation on Actions

Table 6 shows no statistically significant differences in the probability of cooperation in round one $\bar{x}^1$ for supergames where beliefs are elicited. The other regressors are variables that have been considered in similar analysis.

Table 6: Correlated Random Effects Probit
Determinants of Cooperation in Round One

	Finite	Indefinite
Beliefs Are Elicited	0.0654	0.188
	(0.294)	(0.280)
Other Cooperated in Previous Supergame	0.250***	0.624***
	(0.0661)	(0.181)
Supergame	0.0131	0.187***
	(0.0431)	(0.0532)
Length of Previous Supergame		-0.00119
		(0.00807)
Cooperated in Supergame 1	2.571***	2.830***
	(0.754)	(0.649)
Risk Measure	0.0189***	0.00534
	(0.00691)	(0.00663)
Constant	-0.509	-1.461***
	(0.334)	(0.567)
Observations	1778	1126

Standard errors clustered (at the session level) in parentheses. ***1%, **5%, *10% significance.
All variables refer to behavior in Round 1.
Risk Measure is equal to the number of boxes collected in the bomb task.

Table 7: Correlated Random Effects Probit (Marginal Effects)
Determinants of Cooperation in Round One

	Finite	Indefinite
Beliefs Are Elicited	0.00660	0.0193
	(0.0295)	(0.0288)
Other Cooperated in Previous Supergame	0.0252***	0.0642***
	(0.00559)	(0.0188)
Supergame	0.00132	0.0193***
	(0.00439)	(0.00497)
Length of Previous Supergame		-0.000122
		(0.000827)
Cooperated in Supergame 1	0.259***	0.291***
	(0.0790)	(0.0482)
Risk Measure	0.00191***	0.000549
	(0.000615)	(0.000676)
Observations	1778	1126

Standard errors clustered (at the session level) in parentheses. ***1%, **5%, *10% significance.
All variables refer to behavior in Round 1.
Risk Measure is equal to the number of boxes collected in the bomb task.

B.3 Beliefs–Comparative Statics

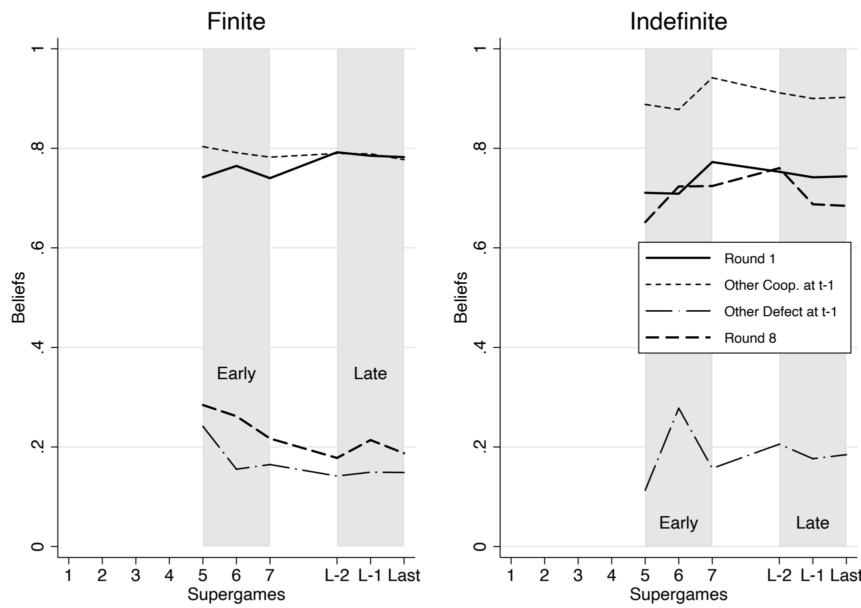

Figure 14: Beliefs Over Supergames

The evolution of beliefs depicted in Figure 14 mirrors the patterns observed for cooperation in Figure 1. $\bar{\mu}^1$ are high in both games. Beliefs are responsive in both games: $\bar{\mu}_i^t(*, a_j^{t-1} = C, *) - \bar{\mu}_i^t(*, a_j^{t-1} = D, *) > 0$. Beliefs $\bar{\mu}^T$ in the last round are low in the Finite game, but are high in the Indefinite game.

B.4 Actions and Beliefs in Rounds One and Seven

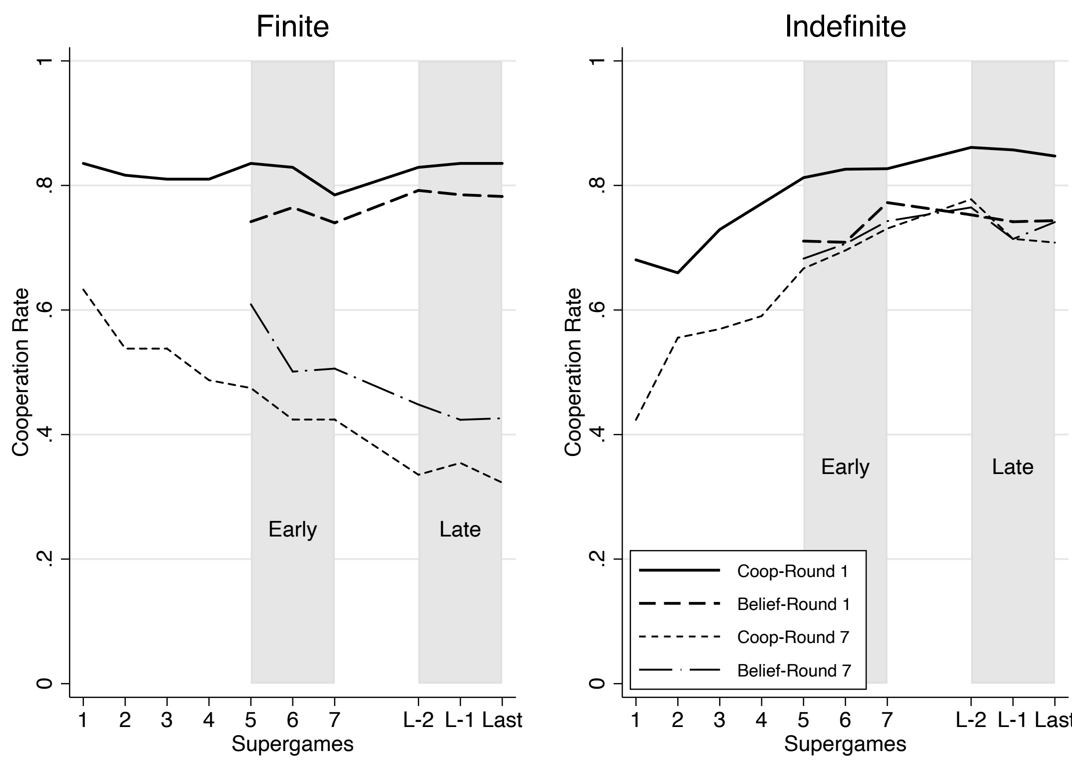

Figure 15: Average Cooperation and Belief

B.5 How Beliefs Relate to Experiences

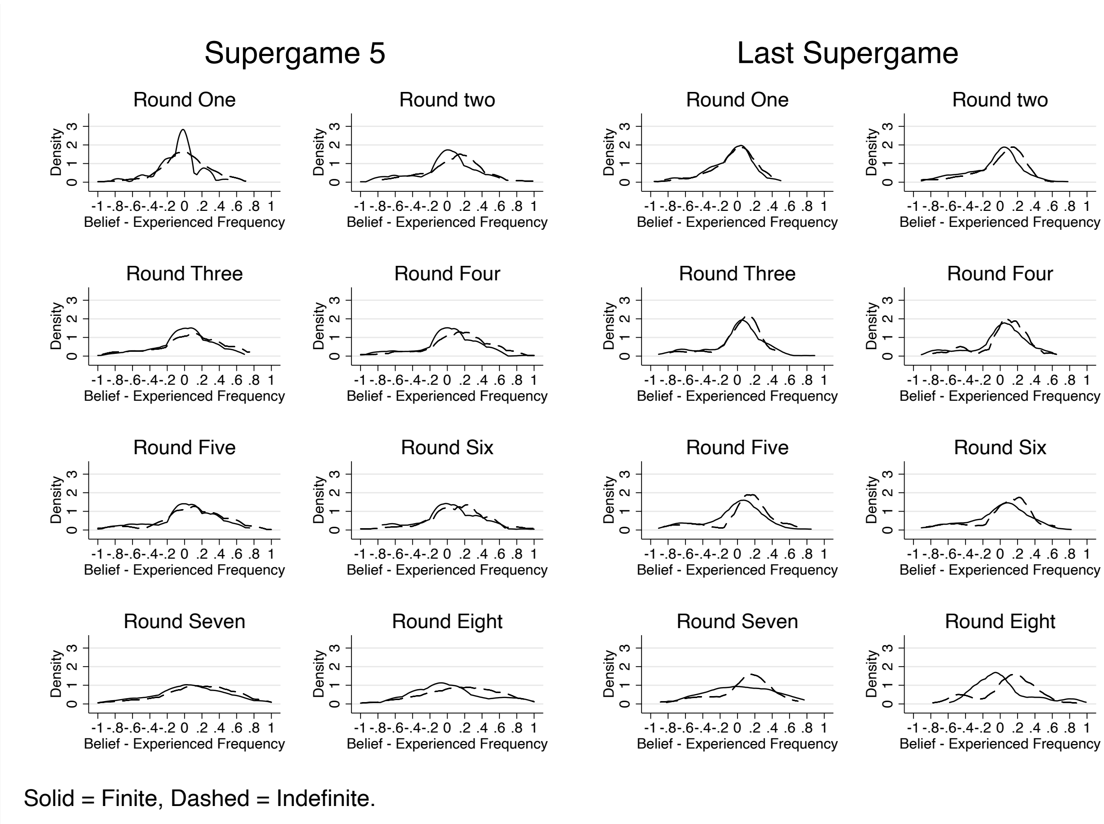

Figure 16: Difference Between Stated Beliefs and Experienced Frequencies of Cooperation by Subject

B.6 Cooperative Path For Early Supergames

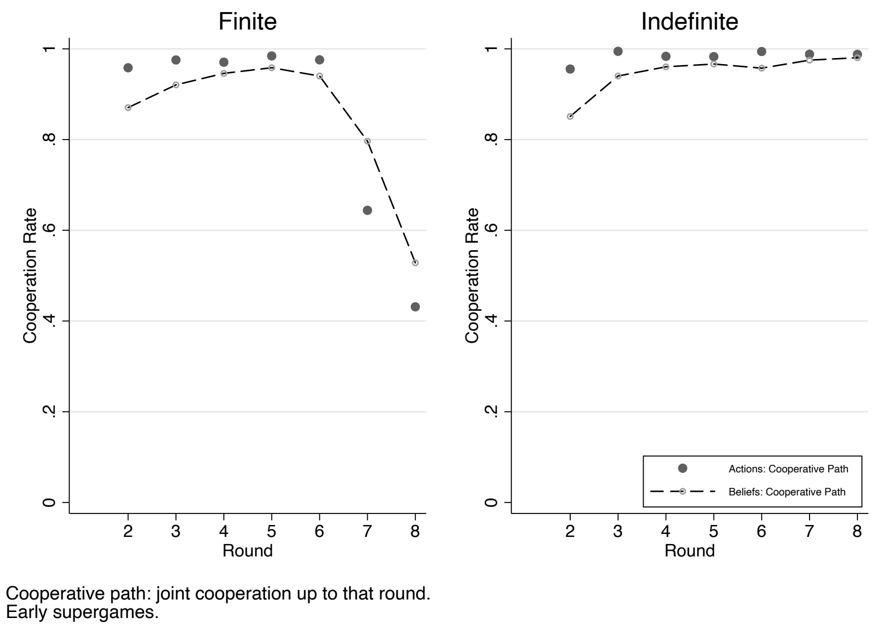

Figure 17: Cooperative Path (First Eight Rounds)

B.7 Accuracy Across Games and Experience

Table 8: Accuracy (numbers are percentages)

Finite

		Early			**Late**		
		Correct	Within		Correct	Within	
		Tercile	10%	5%	Tercile	10%	5%
	Round 1	69	14	8	73	14	7
	Round 2						
	CC	87	60	7	91	60	9
Round 1	CD	63	16	8	67	16	9
Actions	DC	67	11	4	66	7	7
	DD	67	0	0	67	8	8
	Average	80	44	7	83	45	9

Indefinite

		Early			**Late**		
		Correct	Within		Correct	Within	
		Tercile	10%	5%	Tercile	10%	5%
	Round 1	65	13	7	67	10	5
	Round 2						
	CC	86	52	5	91	66	58
Round 1	CD	35	24	12	29	10	2
Actions	DC	65	6	6	56	17	12
	DD	11	0	0	79	0	0
	Average	73	40	6	80	52	45

Round 1 actions are listed own action first, other's action second: i.e. (a_i, a_j).
Average is weighted by the number of observations.
Note: the number of observations following DD is small, with 2% and 5%, for finite and indefinite respectively, of observations for late supergames.

B.8 Relation Between Beliefs and Actions by Treatment

Table 9: Correlated Random Effects Probit (Marginal Effects)
Dependent Variable: Cooperation

	Finite	Indefinite
Belief	0.462***	0.395***
	(0.0176)	(0.0146)
Round	-0.0336***	-0.00238
	(0.00339)	(0.00282)
Coop. in Round 1, Supergames 1-4	0.244***	0.0805***
	(0.0477)	(0.0244)
Coop. in Last Round, Supergames 1-4	0.126***	0.111***
	(0.0164)	(0.0321)
Risk Measure	-0.0000121	0.000105
	(0.000771)	(0.000633)
Observations	3792	3628

Standard errors clustered (at the session level) in parentheses. ***1%, **5%, *10% significance.
Late supergames.
Risk Measure is equal to the number of boxes collected in the bomb task.

We use *round* to make the regressions succinct, but a specification with round indicator variables gives similar estimates.

Table 10: Correlated Random Effects Probit (Marginal Effects)
Dependent Variable: Cooperation in Round One

	Finite	Indefinite
Belief	0.0938***	0.258***
	(0.0272)	(0.0192)
Other Cooperated in Previous Supergame	-0.0382	0.0274
	(0.0379)	(0.0340)
Supergame	0.00143	0.00798
	(0.00925)	(0.00572)
Length of Previous Supergame		-0.00161
		(0.00121)
Cooperated in Supergame 1	0.413***	0.0493***
	(0.0810)	(0.0164)
Risk Measure	0.00163*	-0.000351
	(0.000848)	(0.000561)
Observations	474	378

Standard errors clustered (at the session level) in parentheses. ***1%, **5%, *10% significance.
All variables refer to behavior in Round 1.
Late supergames.
Risk Measure is equal to the number of boxes collected in the bomb task.

Table 11: Description of Strategies Estimated

Name of Strategy	Code	Description
Always Defect	AD	always play D.
Always Cooperate	AC	always play C.
Grim	GRIM	play C until either player plays D, then play D forever.
Tit-For-Tat	TFT	play C unless partner played D last round.
Suspicious Tit-For-Tat	STFT	play D in the first round, then TFT.
Threshold 8	T8	play Grim until round 8 (last round) then switch to AD.
Threshold 7	T7	play Grim until round 7 then switch to AD.
Threshold 6	T6	play Grim until round 6 then switch to AD.
Threshold 5	T5	play Grim until round 5 then switch to AD.
Threshold 4	T4	play Grim until round 4 then switch to AD.
Threshold 3	T3	play Grim until round 3 then switch to AD.
Threshold 2	T2	play C in round 1 then switch to AD.
Lenient Grim 2	GRIM2	play C until 2 consecutive rounds occur in which either player played D, then play D forever.
Tit-For-2 Tats	TF2T	play C unless partner played D in both of the last rounds.
2Tits-For-Tat	2TFT	play C unless partner played D in either of the last 2 rounds.
Lenient Grim 3	GRIM3	play C until 3 consecutive rounds occur in which either player played D, then play D forever.

B.9 Proof of the Cooperativeness Order

When each strategy is denoted by a finite automaton, we assume that an *implementation error* is made in the choice of an action in each state, and not in transition from the current state to the next. We also assume that the errors are independent and identically distributed between the players and across rounds. Denote by $\varepsilon \in [0, \frac{1}{2}]$ the probability of such an error.[65] For the analytical comparison of cooperative levels, we assume that ε is small. In some cases considered below, this implies that we treat ε^2 as negligible. In other cases, however, we need to consider the difference in the order of ε^2 and treat ε^3 as negligible. Let $p = (1-\varepsilon)^2$, $q = \varepsilon(1-\varepsilon)$ and $r = \varepsilon^2$. The normalized stage payoffs with implementation errors are given by

$$\begin{aligned} g_{CC} &= p + q(1+g-\ell), & g_{CD} &= p(-\ell) + q + r(1+g), \\ g_{CD} &= p(1+g) + q + r(-\ell), & g_{DD} &= q(1+g-\ell) + r, \end{aligned}$$

where $g = 1$ and $\ell = 17/12 \approx 1.416$ in our implementation. Define

$$g = \begin{bmatrix} g_{CC} \\ g_{CD} \\ g_{DC} \\ g_{DD} \end{bmatrix}.$$

We consider a Markov process induced by a pair of the same strategy implemented with errors ε. Let Θ be the set of states of this Markov process. For each strategy that can be expressed as an S-state automaton, Θ can have up to $S \times S$ elements. The Markov process is defined over the set $\Delta\Theta$ of distributions over those states. Let $\omega^1 \in \Delta\Theta$ be the row vector representing the initial distribution and $A = (a_{st})_{s,t\in\Theta}$ be the transition matrix: a_{st} is the probability that the next state is t when the current state is s. The distribution ω^2 over round 2 states is given by $\omega^2 = \omega^1 A$, and the distribution ω^t over round t states is given by $\omega^t = \omega^1 A^{t-1}$. With the distribution ω over states, the expected stage payoff to a player is given by ωg. In the case of the finite games, the average payoff over eight rounds can be computed as

$$\frac{1}{8} \sum_{t=1}^{8} \omega^t g = \frac{1}{8} \omega^1 \left(I + A^1 + \cdots + A^7 \right) g. \tag{1}$$

In the case of the indefinite games, the average discounted payoff can be computed as

$$\begin{aligned} (1-\delta) \sum_{t=1}^{\infty} \omega^t \delta^{t-1} g &= (1-\delta)\, \omega^1 \left(I + \delta A^1 + \cdots + \delta^t A^t + \cdots \right) g \\ &= (1-\delta)\, \omega^1 (I - \delta A)^{-1} g, \end{aligned} \tag{2}$$

[65] Hence, $\varepsilon = 1 - \beta$ for the parameter β in SFEM.

where $\delta = 7/8$ in our implementation. If we denote by v_θ the average discounted payoff in the indefinite games along the Markov process with the initial state θ (i.e., the initial distribution ω^1 places probability one on state θ), and by $v = (v_\theta)_{\theta\in\Theta}$ the corresponding column vector, then (2) implies the recursive equation

$$v = (1-\delta)\,(I-\delta A)^{-1}\,g \quad\Leftrightarrow\quad v = (1-\delta)\,g + \delta A v. \tag{3}$$

B.9.1 Indefinite games with small implementation errors

1. TFT and STFT: These strategies have two states 0 and 1. Both strategies play C in state 0, and D in state 1. Because the implementation errors occur independently between the two players, state transitions do not synchronize between them. Accordingly, the Markov process has four states $\Theta = \{(0,0),(0,1),(1,0),(1,1)\}$. The initial distribution is $\omega^1 = (1,0,0,0)$ if both play TFT and $\omega^1 = (0,0,0,1)$ if both play STFT. We hence have $v^{\text{TFT}} = v_{00}$ and $v^{\text{STFT}} = v_{11}$. The transition matrix is given by

$$A = \begin{bmatrix} p & q & q & r \\ q & r & p & q \\ q & p & r & q \\ r & q & q & p \end{bmatrix}.$$

Ignoring the terms of order ε^2, we can write (3) as

$$\begin{bmatrix} v_{00} \\ v_{01} \\ v_{10} \\ v_{11} \end{bmatrix} = (1-\delta)\begin{bmatrix} g_{CC} \\ g_{CD} \\ g_{DC} \\ g_{DD} \end{bmatrix} + \delta \begin{bmatrix} 1-2\varepsilon & \varepsilon & \varepsilon & 0 \\ \varepsilon & 0 & 1-2\varepsilon & \varepsilon \\ \varepsilon & 1-2\varepsilon & 0 & \varepsilon \\ 0 & \varepsilon & \varepsilon & 1-2\varepsilon \end{bmatrix} \begin{bmatrix} v_{00} \\ v_{01} \\ v_{10} \\ v_{11} \end{bmatrix}. \tag{4}$$

It follows from the second and third rows of (4) that

$$\begin{aligned} \begin{bmatrix} v_{01} \\ v_{10} \end{bmatrix} &= (1-\delta)\begin{bmatrix} g_{CD} \\ g_{DC} \end{bmatrix} + \delta \begin{bmatrix} v_{10} \\ v_{01} \end{bmatrix} + \delta\varepsilon \begin{bmatrix} v_{00}+v_{11}-2v_{10} \\ v_{00}+v_{11}-2v_{01} \end{bmatrix} \\ &= (1-\delta)\begin{bmatrix} g_{CD} \\ g_{DC} \end{bmatrix} + \delta \begin{bmatrix} v_{10} \\ v_{01} \end{bmatrix} + O(\varepsilon), \end{aligned}$$

where $O(\varepsilon)$ is the term of order ε. Hence,

$$\begin{bmatrix} 1 & -\delta \\ -\delta & 1 \end{bmatrix} \begin{bmatrix} v_{01} \\ v_{10} \end{bmatrix} = (1-\delta)\begin{bmatrix} g_{CD} \\ g_{DC} \end{bmatrix} + O(\varepsilon).$$

Solving this, we get

$$\begin{bmatrix} v_{01} \\ v_{10} \end{bmatrix} = \frac{1}{1+\delta}\begin{bmatrix} 1 & \delta \\ \delta & 1 \end{bmatrix} \begin{bmatrix} g_{CD} \\ g_{DC} \end{bmatrix} + O(\varepsilon).$$

Substituting this into the first and fourth rows of (4), we obtain

$$\begin{aligned}\begin{bmatrix} v_{00} \\ v_{11} \end{bmatrix} &= (1-\delta) \begin{bmatrix} g_{CC} \\ g_{DD} \end{bmatrix} + \delta(1-2\varepsilon) \begin{bmatrix} v_{00} \\ v_{11} \end{bmatrix} + \delta\varepsilon(1+\delta) \begin{bmatrix} v_{01}+v_{10} \\ v_{01}+v_{10} \end{bmatrix} \\ &= (1-\delta) \begin{bmatrix} g_{CC} \\ g_{DD} \end{bmatrix} + \delta(1-2\varepsilon) \begin{bmatrix} v_{00} \\ v_{11} \end{bmatrix} + \delta\varepsilon \begin{bmatrix} g_{CD}+g_{DC} \\ g_{CD}+g_{DC} \end{bmatrix} + O(\varepsilon^2).\end{aligned}$$

This can be rewritten as

$$\begin{aligned}&\begin{bmatrix} 1-\delta+2\delta\varepsilon & 0 \\ 0 & 1-\delta+2\delta\varepsilon \end{bmatrix} \begin{bmatrix} v_{00} \\ v_{11} \end{bmatrix} \\ &= (1-\delta) \begin{bmatrix} g_{CC} \\ g_{DD} \end{bmatrix} + \delta\varepsilon \begin{bmatrix} g_{CD}+g_{DC} \\ g_{CD}+g_{DC} \end{bmatrix} + O(\varepsilon^2).\end{aligned}$$

Ignoring the terms involving ε^2, we hence obtain

$$\begin{bmatrix} v^{\text{TFT}} \\ v^{\text{STFT}} \end{bmatrix} = \begin{bmatrix} v_{00} \\ v_{11} \end{bmatrix} = \frac{1}{1-\delta+2\delta\varepsilon} \begin{bmatrix} (1-\delta)\, g_{CC} + \delta\varepsilon(g_{CD}+g_{DC}) \\ (1-\delta)\, g_{DD} + \delta\varepsilon(g_{CD}+g_{DC}) \end{bmatrix}.$$

2. Grim: The strategy has two states 0 and 1 where it chooses C and D, respectively. State transitions are synchronized between the two players when they both play Grim so that the Markov process has only two states $\Theta = \{(0,0),(1,1)\}$. We have $\omega^1 = (1,0)$ so that $v^{\text{Grim}} = v_{00}$. The transition matrix is given by

$$A = \begin{bmatrix} p & 1-p \\ 0 & 1 \end{bmatrix}.$$

Ignoring the terms of order ε^2, we can write (3) as

$$\begin{bmatrix} v_{00} \\ v_{11} \end{bmatrix} = (1-\delta) \begin{bmatrix} g_{CC} \\ g_{DD} \end{bmatrix} + \delta \begin{bmatrix} 1-2\varepsilon & 2\varepsilon \\ 0 & 1 \end{bmatrix} \begin{bmatrix} v_{00} \\ v_{11} \end{bmatrix}$$

This yields

$$v^{\text{Grim}} = v_{00} = \frac{(1-\delta)g_{CC} + 2\delta\varepsilon g_{DD}}{1-\delta+2\delta\varepsilon}.$$

3. Grim2: The strategy has three states 0, 1 and 2, where it chooses C, C, and D, respectively. State transitions are synchronized between the two players so that the Markov process has three states $\Theta = \{(0,0),(1,1),(2,2)\}$. We have $\omega^1 = (1,0,0)$ so that $v^{\text{Grim2}} = v_{00}$. The transition matrix is given by

$$A = \begin{bmatrix} p & 1-p & 0 \\ p & 0 & 1-p \\ 0 & 0 & 1 \end{bmatrix}.$$

We can write (3) as

$$\begin{bmatrix} v_{00} \\ v_{11} \\ v_{22} \end{bmatrix} = (1-\delta) \begin{bmatrix} g_{CC} \\ g_{CC} \\ g_{DD} \end{bmatrix} + \delta \begin{bmatrix} (1-\varepsilon)^2 & \varepsilon(2-\varepsilon) & 0 \\ (1-\varepsilon)^2 & 0 & \varepsilon(2-\varepsilon) \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} v_{00} \\ v_{11} \\ v_{22} \end{bmatrix}.$$

Solving this, we obtain

$$v^{\text{Grim2}} = v_{00} = \frac{(1-\delta)\{1+\delta\varepsilon(2-\varepsilon)\}g_{CC} + 4\delta^2\varepsilon^2 g_{DD}}{(1-\delta)\{1+\delta\varepsilon(2-\varepsilon)\} + 4\delta^2\varepsilon^2}.$$

4. TF2T: The strategy has three states 0, 1 and 2, where the action choices are C, C, and D, respectively. Since state transitions are not synchronized, the Markov process has $3 \times 3 = 9$ states $\Theta = \{(0,0), \ldots, (2,2)\}$. We have $\omega^1 = (1, 0, \ldots, 0)$ so that $v^{\text{TF2T}} = v_{00}$. The transition matrix is given by

$$A = \begin{bmatrix} p & q & 0 & q & r & 0 & 0 & 0 & 0 \\ p & 0 & q & q & 0 & r & 0 & 0 & 0 \\ q & 0 & r & p & 0 & q & 0 & 0 & 0 \\ p & q & 0 & 0 & 0 & 0 & q & r & 0 \\ p & 0 & q & 0 & 0 & 0 & q & 0 & r \\ q & 0 & r & 0 & 0 & 0 & p & 0 & q \\ q & p & 0 & 0 & 0 & 0 & r & q & 0 \\ q & 0 & p & 0 & 0 & 0 & r & 0 & q \\ r & 0 & q & 0 & 0 & 0 & q & 0 & p \end{bmatrix}.$$

Using (3), we have

$$\begin{aligned} v_{11} &= (1-\delta)g_{CC} + \delta v_{00} + O(\varepsilon) \\ v_{02} &= (1-\delta)g_{CD} + \delta v_{10} + O(\varepsilon) \\ v_{20} &= (1-\delta)g_{DC} + \delta v_{01} + O(\varepsilon). \end{aligned} \tag{5}$$

Substituting these into the recursive equations for v_{01} and v_{10}, we obtain

$$\begin{aligned} \begin{bmatrix} v_{01} \\ v_{10} \end{bmatrix} = (1-\delta) \begin{bmatrix} g_{CC} \\ g_{CC} \end{bmatrix} + \delta(1-2\varepsilon) \begin{bmatrix} v_{00} \\ v_{00} \end{bmatrix} + \delta(1-\delta)\varepsilon \begin{bmatrix} g_{CD} \\ g_{DC} \end{bmatrix} \\ + \delta\varepsilon \begin{bmatrix} 0 & 1+\delta \\ 1+\delta & 0 \end{bmatrix} \begin{bmatrix} v_{01} \\ v_{10} \end{bmatrix} + O(\varepsilon^2), \end{aligned}$$

which yields

$$\begin{aligned} \begin{bmatrix} v_{01} \\ v_{10} \end{bmatrix} = \frac{1-\delta}{1-\delta^2\varepsilon^2(1+\delta)^2} \begin{bmatrix} 1 & \delta\varepsilon(1+\delta) \\ \delta\varepsilon(1+\delta) & 1 \end{bmatrix} \begin{bmatrix} g_{CC} + \delta\varepsilon g_{CD} \\ g_{CC} + \delta\varepsilon g_{DC} \end{bmatrix} \\ + \frac{\delta(1-2\varepsilon)v_{00}}{1-\delta^2\varepsilon^2(1+\delta)^2} \begin{bmatrix} 1 & \delta\varepsilon(1+\delta) \\ \delta\varepsilon(1+\delta) & 1 \end{bmatrix} \begin{bmatrix} 1 \\ 1 \end{bmatrix} + O(\varepsilon^2). \end{aligned}$$

It then follows that

$$
\begin{aligned}
v_{01} + v_{10} &= \frac{(1-\delta)\{1+\delta\varepsilon(1+\delta)\}}{1-\delta^2\varepsilon^2(1+\delta)^2}\ \{2g_{CC} + \delta\varepsilon(g_{CD}+g_{DC})\} \\
&\quad + \frac{2\delta(1-2\varepsilon)\{1+\delta\varepsilon(1+\delta)\}}{1-\delta^2\varepsilon^2(1+\delta)^2}\, v_{00} + O(\varepsilon^2) \\
&= \frac{(1-\delta)}{1-\delta\varepsilon(1+\delta)}\ \{2g_{CC} + \delta\varepsilon(g_{CD}+g_{DC})\} \\
&\quad + \frac{2\delta(1-2\varepsilon)}{1-\delta\varepsilon(1+\delta)}\, v_{00} + O(\varepsilon^2).
\end{aligned}
\tag{6}
$$

On the other hand, the recursive equation for v_{00} yields

$$
v_{00} = \frac{(1-\delta)g_{CC} + \delta\varepsilon(1-\varepsilon)(v_{01}+v_{10}) + \delta\varepsilon^2 v_{11}}{1-\delta(1-\varepsilon)^2}. \tag{7}
$$

Substituting (5) and (6) into (7) and ignoring the terms of order ε^3, we obtain

$$
v^{\mathrm{TF2T}} = v_{00} = \frac{\{1+\delta(1-\delta)\varepsilon - \delta\varepsilon^2\}g_{CC} + \delta^2\varepsilon^2(g_{CD}+g_{DC})}{1+\delta(1-\delta)\varepsilon - \delta(1-2\delta)\varepsilon^2}.
$$

As for the strategies AC, AD, and T6-T8, it can be readily verified that their cooperativeness is given as follows.

5. AD: $v^{\mathrm{AD}} = g_{DD}$.
6. AC: $v^{\mathrm{AC}} = g_{CC}$.
7. T8: $v^{\mathrm{T8}} = (1-\delta^7)\, g_{CC} + \delta^7\, g_{DD} + O(\varepsilon)$.
8. T7: $v^{\mathrm{T7}} = (1-\delta^6)\, g_{CC} + \delta^6\, g_{DD} + O(\varepsilon)$.
9. T6: $v^{\mathrm{T6}} = (1-\delta^5)\, g_{CC} + \delta^5\, g_{DD} + O(\varepsilon)$.

Combining the above cases, we can rank the ten strategies from the least cooperative to the most cooperative in the indefinite games as follows:

$$
\begin{aligned}
&\mathrm{AD} \ll \mathrm{STFT} \lll \mathrm{T6} \lll \mathrm{T7} \lll \mathrm{T8} \\
&\lll \mathrm{Grim} \ll \mathrm{TFT} \ll \mathrm{Grim2} < \mathrm{TF2T} < \mathrm{AC},
\end{aligned}
$$

where $\lll$, $\ll$ and $<$ represent domination in the orders of $\varepsilon^0(=1)$, ε, and ε^2, respectively.

B.9.2 General implementation errors

When the probability $\varepsilon \in [0, \frac{1}{2}]$ of implementation errors is not necessarily small, the cooperativeness of the strategies TFT, STFT, Grim, Grim2, and TF2T can be computed numerically using (1) for the finite games and by (2) for the indefinite games, whereas the cooperativeness of AC and AD equals g_{CC} and g_{DD}, respectively, as above. Consider now the strategy Tk ($k = 6$, 7, 8). In the indefinite games, its cooperativeness can be computed as

$$v^{\mathrm{T}k} = (1-\delta)\,\frac{1-(\delta p)^{k-1}}{1-\delta p}\,g_{CC} + \delta\left\{(1-p)\frac{1-(\delta p)^{k-2}}{1-\delta p} + (\delta p)^{k-2}\right\}\,g_{DD}.$$

In the finite games, suppose that $t < k$ and let v_t denote the sum of stage payoffs in rounds $t, t+1, \ldots, 8$ when Tk still specifies action C in round t. We have the following recursive equations:

$$\begin{aligned}
v_{k-1} &= g_{CC} + (9-k)g_{DD}, \\
v_{k-2} &= g_{CC} + pv_{k-1} + (1-p)(10-k)g_{DD}, \\
&\ \,\vdots \\
v_2 &= g_{CC} + pv_3 + (1-p)\cdot 6g_{DD}, \\
v_1 &= g_{CC} + pv_2 + (1-p)\cdot 7g_{DD}.
\end{aligned}$$

The cooperativeness of Tk then equals $v^{\mathrm{T}k} = \frac{v_1}{8}$.

B.10 Complete Estimation Results

Table 12: 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	$\tilde{\beta}$
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)		
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)		
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)		
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)		
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)		
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)		
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)		
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)		
GRIM2	0.02	0.01	-	-	-	-	-	-	-	-	-	-	-	-
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)		
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 13: 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	$\tilde{\beta}$
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)		
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)		
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)		
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)		
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	-	-	-	-	-	-	-	-	-	-	-	-
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)		
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)		
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 14: 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	$\tilde{\beta}$
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)		
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	-	-	-	-	-	-	-	-	-	-	-	-
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)		
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)		
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)		
T6	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-
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)		
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)		
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 15: 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	$\tilde{\beta}$
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)		
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)		
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)		
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)		
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	-	-	-	-	-	-	-	-	-	-	-	-
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)		
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)		
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.

Table 16: Estimates for the Finite Game on Late Supergames

	Share		Estimated Beliefs - $\tilde{p}$																	
Type	SFEM	Typing	AD	AC	GRIM	TFT	STFT	T8	T7	T6	T5	T4	T3	T2	GRIM2	TF2T	2TFT	GRIM3	ν	$\tilde{\beta}$
AD	0.12	0.12	0.12	0.00	*[0.00]*	0.00	0.13	*[0.00]*	0.00	*[0.00]*	*[0.00]*	*[0.00]*	*[0.00]*	0.00	0.00	0.00	0.75	0.00	0.06	1.00
AC	0.02	0.01	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
GRIM	0.07	0.02	0.06	0.06	0.11	0.20	0.27	0.23	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.06	*[0.00]*	0.02	0.07	1.00
TFT	0.09	0.12	0.11	0.00	0.00	0.55	0.03	0.30	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.05	1.00
STFT	0.02	0.02	0.00	0.00	*[0.00]*	0.81	0.00	*[0.00]*	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.19	0.00	0.00	0.14	1.00
T8	0.22	0.20	0.05	0.00	0.00	0.00	0.00	0.52	0.00	0.00	0.00	0.00	0.00	0.00	0.15	0.15	0.00	0.13	0.04	1.00
T7	0.30	0.35	0.00	0.00	0.00	0.36	0.00	0.23	0.41	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.04	1.00
T6	0.08	0.08	0.00	0.00	*[0.00]*	0.00	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.03	1.00
T5	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
T4	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
T3	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
T2	0.01	0.01	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
GRIM2	0.01	0.01	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
TF2T	0.03	0.04	0.00	0.17	0.83	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.03	1.00
2TFT	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
GRIM3	0.03	0.03	0.00	0.00	*[0.00]*	0.78	0.00	0.00	0.00	0.00	0.22	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.07	1.00
ALL			0.04	0.01	0.04	0.22	0.04	0.23	0.21	0.00	0.01	0.00	0.00	0.00	0.03	0.04	0.09	0.09		

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

Table 17: Estimates for the Indefinite Game on Late Supergames

	Share		Estimated Beliefs - $\tilde{p}$																	
Type	SFEM	Typing	AD	AC	GRIM	TFT	STFT	T8	T7	T6	T5	T4	T3	T2	GRIM2	TF2T	2TFT	GRIM3	ν	$\tilde{\beta}$
AD	0.09	0.10	0.90	0.01	0.07	0.01	0.00	*[0.00]*	*[0.00]*	*[0.00]*	*[0.00]*	*[0.00]*	*[0.00]*	*[0.00]*	0.01	0.00	0.00	0.01	0.04	1.00
AC	0.10	0.10	0.00	0.85	0.00	0.15	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.05	1.00
GRIM	0.15	0.07	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.06	1.00
TFT	0.34	0.58	0.08	0.12	0.08	0.28	0.04	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.26	0.14	0.00	0.01	1.00
STFT	0.04	0.04	0.48	0.00	0.15	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.21	0.14	0.00	0.00	0.00	0.07	1.00
T8	0.01	0.01	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
T7	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
T6	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
T5	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
T4	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
T3	0.02	0.03	0.00	0.00	0.16	0.30	0.00	*[0.00]*	*[0.00]*	*[0.00]*	0.00	0.07	0.07	0.00	0.14	0.03	0.23	0.00	0.08	1.00
T2	0.00	0.00	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
GRIM2	0.07	0.02	0.00	0.24	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.52	0.24	*[0.00]*	0.00	0.05	1.00
TF2T	0.09	0.03	0.00	0.97	0.00	0.02	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	1.00
2TFT	0.05	0.01	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
GRIM3	0.06	0.02	0.00	0.01	0.24	0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.72	0.01	*[0.00]*	0.01	0.01	1.00
ALL			0.14	0.24	0.21	0.13	0.02	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.09	0.11	0.06	0.06		

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

Table 18: Best Response Analysis

Finite						Indefinite					
	Share		Best Response				Share		Best Response		
Type	SFEM	Typing	BRS	R_s	R_o	Type	SFEM	Typing	BRS	R_s	R_o
T7	0.30	0.35	T7	1	0.97	TFT	0.34	0.58	TF2T/GRIM2	0.99	0.93
T8	0.22	0.20	T7	0.89	0.87	GRIM	0.15	0.07	GRIM	1	0.92
AD	0.12	0.12	T8	0.23	0.6	AC	0.10	0.10	AD	0.78	0.74
TFT	0.09	0.12	T8	0.87	0.77	AD	0.09	0.10	AD	1	0.76
T6	0.08	0.08	T6	1	1	TF2T	0.09	0.03	STFT	0.96	0.95
GRIM	0.07	0.02	T7	0.84	0.82	GRIM2	0.07	0.02	STFT	0.89	1
Other	0.12	0.11	T6			Other	0.16	0.10	TFT		
All			T7			All			TFT		

Estimation on late supergames out of 16 strategies: AD, AC, Grim, Grim2, Grim3, TFT, TF2T, 2TFT, STFT, T2-T8.

Rows represent top 6 played strategies; BRS: Best Response strategy given beliefs.

In Finite games the best response strategy to the actual distribution (SFEM) is T6; in Indefinite games it is GRIM2.

R_s: Expected payoff from strategy/Best response payoff given beliefs.

R_o: Expected payoff from strategy/Best response payoff given actual distribution (SFEM).

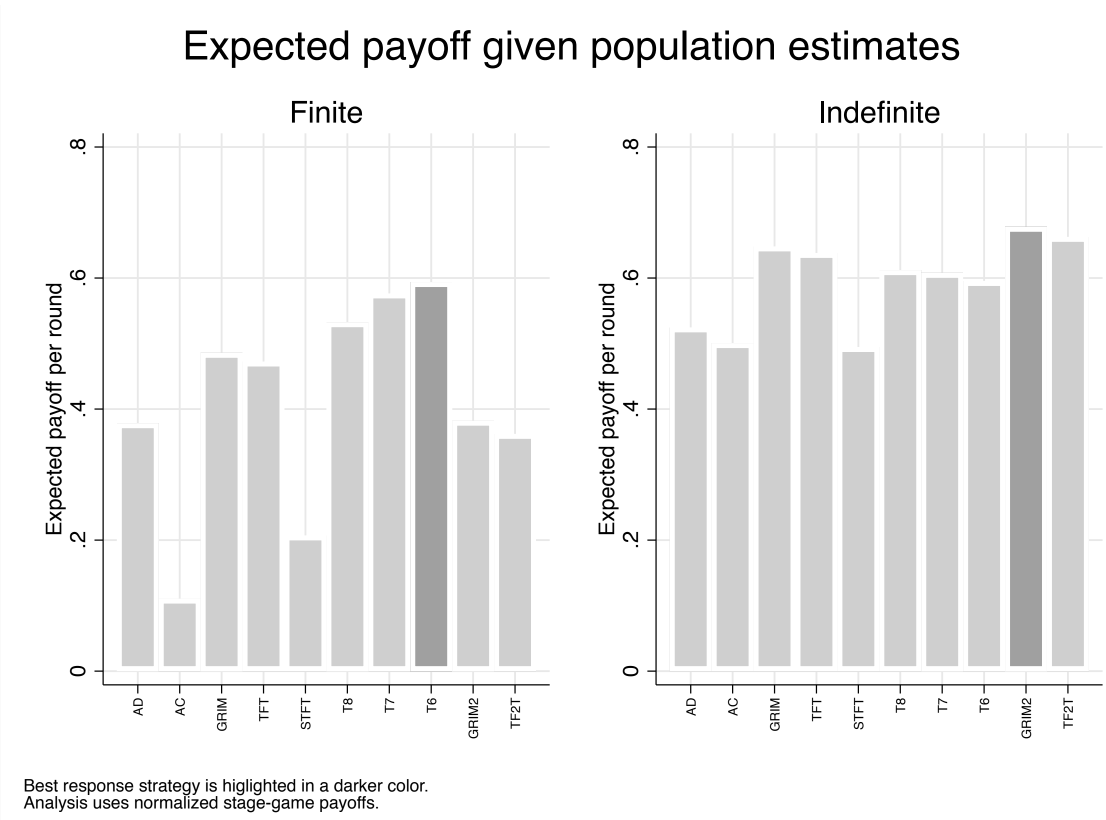

Figure 18: Normalized Expected Payoff by Type Given Strategy Distribution

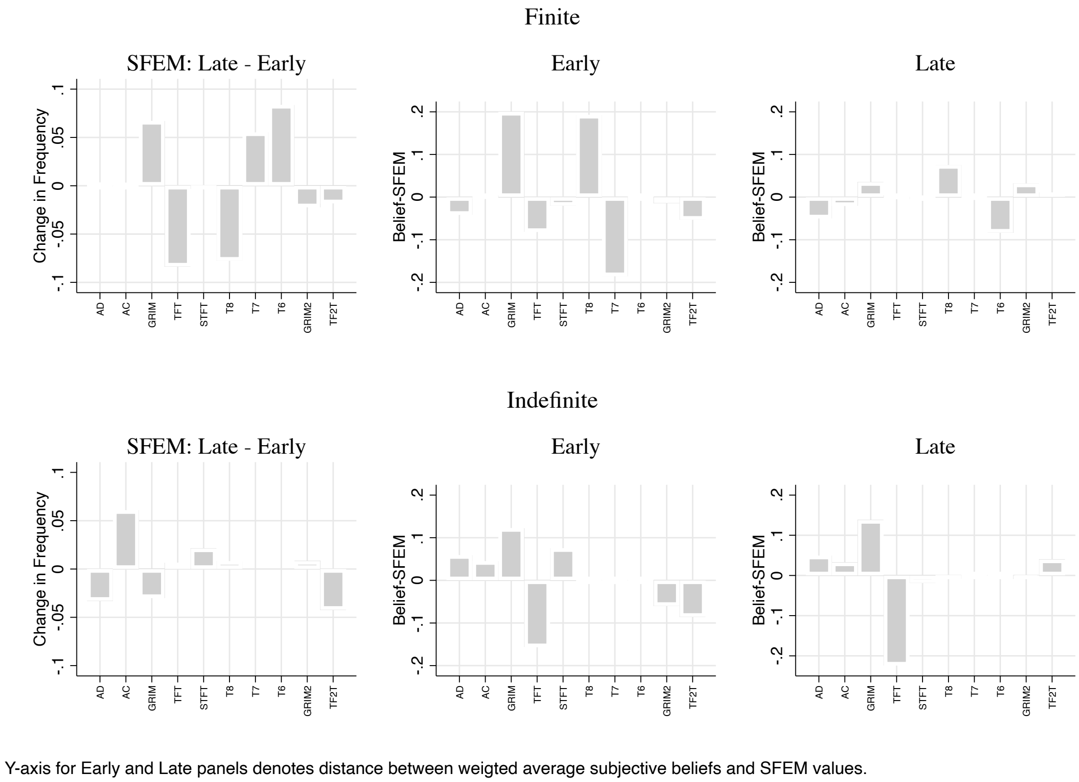

Figure 19: Strategy Changes and Belief Accuracy

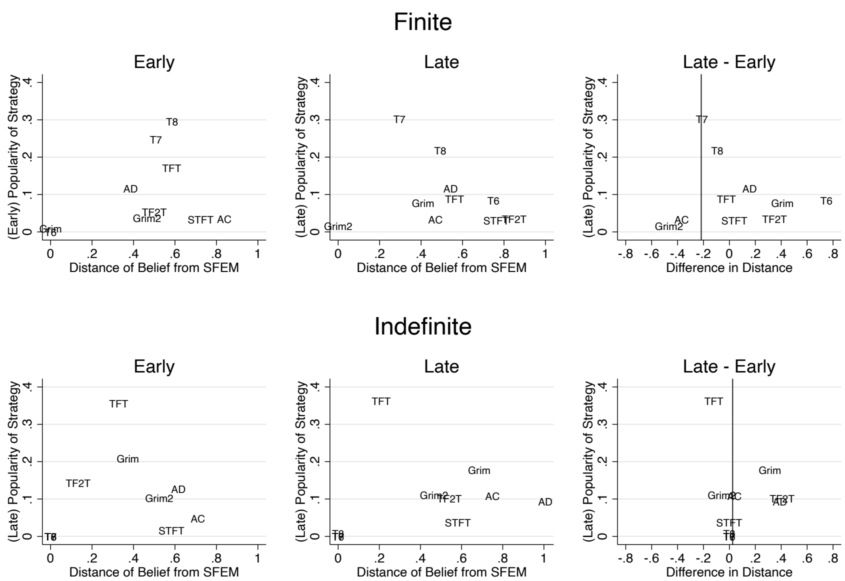

Vertical lines indicate the average weighted by popularity of strategies.

Figure 20: Change in Accuracy

Table 19: Type Evolution: Finite

	T7	T8	AD	TFT	T6	Grim	TF2T	AC	STFT	Grim2
Early Types	32	57	20	24	0	0	12	3	6	4
Number that Change	8	39	4	15			11	3	5	4
No. 1 Change (%)	T6 (50%)	T7 (62%)		T8 (47%)			TFT (27%)		T7 (40%)	T8 (75%)
No. 2 Change (%)		T6 (15%)		T7 (20%)						TF2T (25%)

Sorted by late frequency.
Last two rows provided if no. 1 and 2 are unique.

Table 20: Type Evolution: Indefinite

	TFT	Grim	Grim2	AC	TF2T	AD	STFT	T8	T6	T7
Early Types	86	13	7	3	14	19	2			
Number that Change	25	11	5	3	14	10	2			
No. 1 Change (%)	TF2T (44%)	TFT (64%)	TFT (80%)		TFT (79%)	STFT (40%)	AC (100%)			
No. 2 Change (%)	Grim (32%)	AD (18%)	Grim (20%)							

Sorted by late frequency.
Last two rows provided if no. 1 and 2 are unique.

B.11 Belief-Estimation Method

In this section, we document properties of our belief-estimation method using simulations.

First, we show that the method recovers the correct beliefs in a simple stylized example of a population that consists of only AD (25%) , Grim (40%) and TFT types (35%). We simulate data—including both actions and round-by-round beliefs—based on the model of belief formation described in the paper assuming the following supergame beliefs for the different types. AD types believe others are playing AD with 40% probability, Grim with 10% probability, and TFT with 50% probability. Grim types believe others are playing AD with 10% probability, Grim with 30% probability, and TFT with 60% probability. TFT believe others are playing AD with 20% probability, Grim with 50% probability, and TFT with 30% probability.[66]

Figure 21 plots how well these input parameters are recovered by our belief estimation method from the simulated data. Note that this involves all three steps of our method: applying SFEM to the simulated data, estimating strategy type of each simulated subject taking population level SFEM estimates as a prior, and, finally, for each strategy type estimating beliefs over strategies given simulated round-by-round beliefs. We contrast results from experiments with two sessions, four sessions and eight sessions.

Next we consider a similar exercise but for conditions similar to the ones in our data set. Namely, the data generating process is assumed to correspond to the one we report in Tables 16 and 17. The sample size is assumed to be the same as the one we have collected in the experiment. Figures 22 and 23 show that the input parameters are recovered quite well for the most common types.[67]

[66]We set $\beta = 0.9365$, the average estimated value for this parameter in the experiment (using values from the Finite and Indefinite games). We set $\tilde{\beta} = 1.00$ and $\nu = 0.05$ which are the median estimated values for these parameters from the experiment (including all types in the Finite and Indefinite games).

[67]One notable exception is the supergame beliefs of the AC in the Indefinite game, which are not recovered as well as other types. However, it is useful to note the nature of the discrepancy in this case: input values are such that the AC type puts 80% probability on others playing AC; the recovered values are such that some of this weight is shifted to TF2T. Thus, the discrepancy between the input and output values are among the most cooperative two strategies.

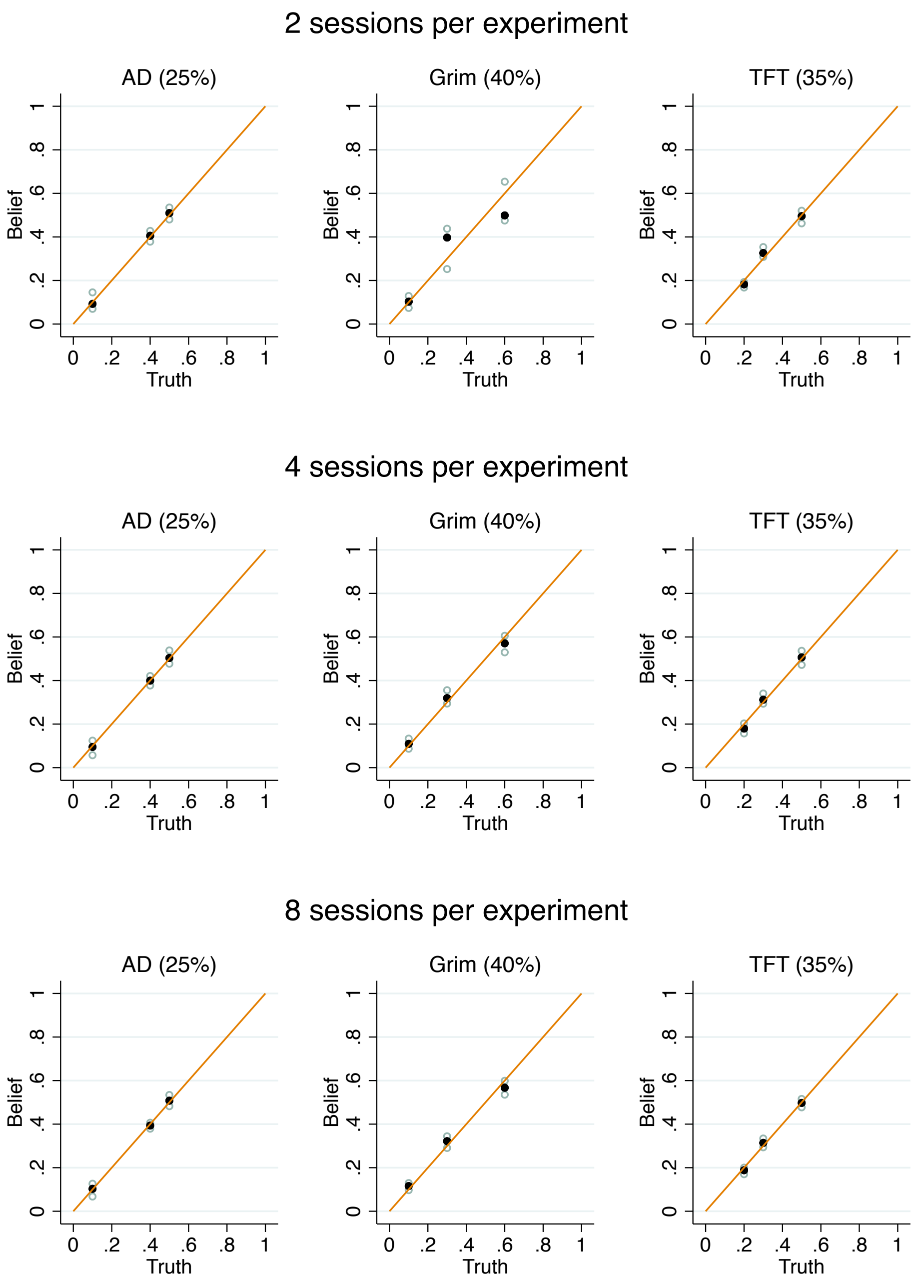

Estimation results from 100 simulated experiments with 18 subjects in each session.
Truth refers to input values.
Solid dots represent median estimate, hollow bubbles represent 25th and 75th percentile estimates.
Input values for all other parameters (not depicted in graphs) correspond to median values from belief estimates reported in the paper.

Figure 21: Estimation results using simulations

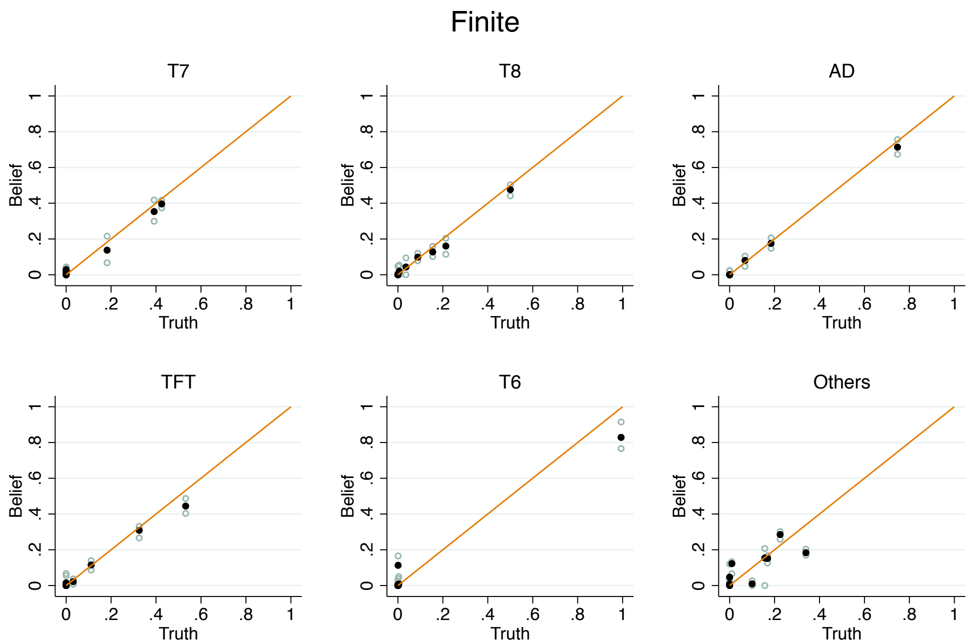

Truth refers to input values.
Solid dots represent median estimate, hollow bubbles represent 25th and 75th percentile estimates from 150 simulations.
We report simulation results for which identification matches original experiment (96% of data).

Figure 22: Estimation results using simulations

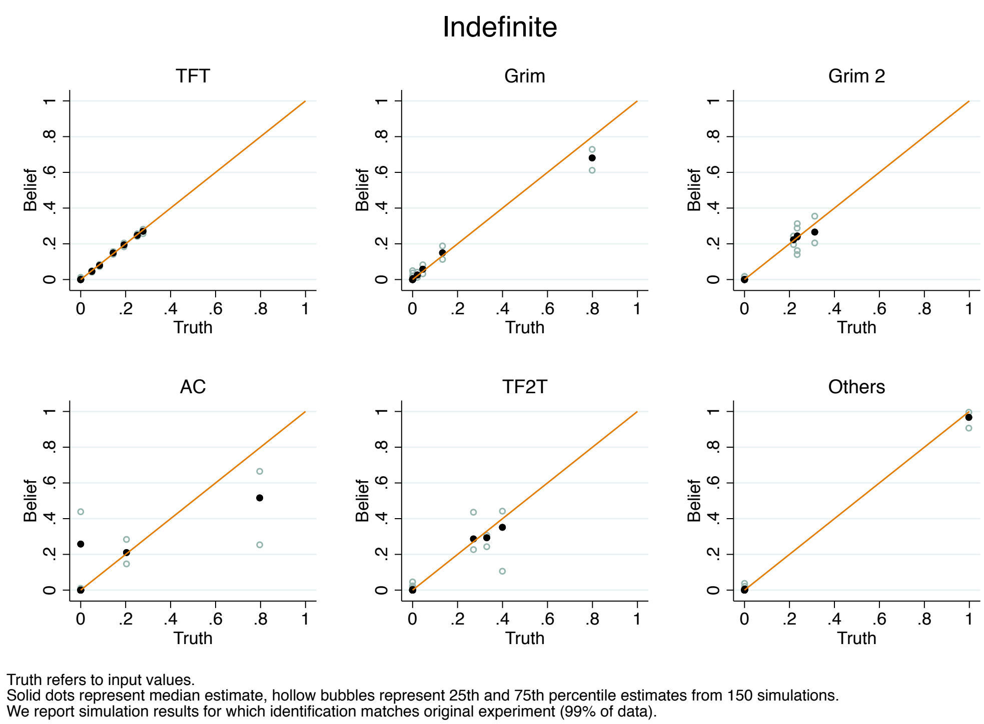

Figure 23: Estimation results using simulations

C Instructions

F8

INSTRUCTIONS

You are about to participate in an experiment on decision-making. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance. Please turn off cell phones and similar devices now. Please do not talk or in any way try to communicate with other participants.

We will start with a brief instruction period. If you have any questions during this period, raise your hand and your question will be answered so everyone can hear.

This experiment has three parts; these instructions are for the first part. Once this part is over, instructions for the next part will be given to you. Your decisions in this part have no influence on the other parts.

General Instructions

1. In this experiment you will be repeatedly matched with a randomly selected person in the room. During each match, you will be asked to make decisions over a sequence of rounds.

2. The points you can obtain in each round of a match depend on your choice and the choice of the person you are paired with. The table below represents all the possible outcomes:

Your Choice	**Other's Choice**	
	1	***2***
1	51, *51*	22 *63*
2	63, *22*	39, *39*

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

   That is, in each round of a match, if:
   - (1,1): Your choice is 1 and the other's choice is 1, you each make 51.
   - (1,2): Your choice is 1 and the other's choice is 2, you make 22 while the other makes 63.
   - (2,1): Your choice is 2 and the other's choice is 1, you make 63 while the other makes 22.
   - (2,2): Your choice is 2 and the other's choice is 2, you each make 39.

3. At the end of each round, you will see your choice (**1** or **2**) and the choice of the person you were paired with (**1** or **2**).

4. Each match will last for 8 rounds.

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

6. Each part of the experiment will generate points that count towards your final payoff. In this part, one match will be randomly selected to count towards your final payoff. Points earned in this match will be converted to dollars at a rate of 3 cents per point. You will receive an additional $8 show up fee for your participation. You will only be informed of your payoffs at the end of the experiment.

7. This part will last for four matches.

Are there any questions?

Before we start, let me remind you that:

- Each match will last for 8 rounds. You will interact with the same person for the entire match.

- Your choice and the choice of the person you are paired with will be shown to both of you at the end of the round.

- Points obtained in each round depend on these choices.

- After a match is finished, you will be randomly paired with someone for a new match.

General Instructions for Part 2

The basic structure of this part is very similar to part 1. How the match proceeds and how you are paired with others will remain the same.

However, in this part, you will have one more task. In each round of a match, after you make a choice, we will ask you to submit your belief about the choice of the person you are paired with.

To indicate your beliefs, you will use a slider. Where you move the slider will represent your best assessment of the likelihood (expressed as chance out of 100) that the person you are paired with chose **1** or **2.**

Two different matches from this part will be randomly selected to count towards payment. For one of these, you will receive the points associated with your choices as in part 1. For the other, the computer will randomly choose one round from that match for payment for beliefs. The belief that you report in that round will determine your chance of winning a prize of 50 points.

To determine your payment, the computer will randomly draw two numbers. For each draw, all numbers between 0 and 100 (including decimal numbers) are equally likely to be selected. Draws are independent in the sense that the outcome of the first draw in no way affects the outcome of the second draw.

If the person you are paired with chose **1** in that round and the number you indicated as the likelihood that the other chose **1** is larger than either of the two draws, you will win the prize.

If the person you are paired with chose **2** in that round and the number you indicated as the likelihood that the other chose **2** is larger than either of the two draws, you will win the prize.

The rules that determine your chance of winning this prize were purposefully designed so that you have the greatest chance of winning the prize when you answer the question with your true assessment on how likely the person you are paired with chose **1** or **2**.

The first match to end after 60 minutes of play (including the first part) will mark the end of the experiment.

General Instructions for Part 3

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

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

There is also a Start button—please do not click on this button until we finish reading the instructions. Once the Start button is clicked, the experiment begins. Every two seconds, a box will be collected, beginning with Box #1 (top left) and ending with Box #100 (bottom right).

You earn 3 cents for every box that is collected. Once collected, the box changes from dark grey to light grey, and your earnings are updated accordingly. At any moment, on the information box, you can see the number of boxes collected so far and the amount earned up to that point.

Such earnings are only ***potential***, however, because behind one of these boxes a bomb is hidden that destroys everything that has been collected in this part of the experiment. You do not know the location of the bomb. Moreover, even if you collect the bomb, you will not know it until the end of the experiment. Your task is to choose when to stop the collecting process. You stop the process by hitting ‘Stop’ at any time.

Payoffs: If at the moment you hit ‘Stop’ none of the boxes you have collected contain the bomb, you will receive the amount of money you have accumulated. If at the moment you hit ‘Stop’ you happen to have collected the box with the bomb, then you will earn $0. Remember that you will not be told if a box that you have collected has or does not have the bomb until after you hit the ‘Stop’ button. So the earnings you see on the screen are only potential earnings, and you will earn those earnings only if none of the boxes you have collected had the bomb.

Location of the Bomb: The interface will randomly choose a number between 1 and 100. All numbers are equally likely. The interface will then place the bomb in the box with the randomly chosen number.