Beliefs in Repeated Games*

Masaki Aoyagi
Osaka University

Guillaume R. Fréchette
NYU

Sevgi Yuksel
UCSB

July 9, 2020

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 that subjects' beliefs elicited in each round about the other player's action are remarkably accurate despite some important systematic deviations corresponding to early pessimism in indefinite games and late optimism in finite games. The data reveals a close link between beliefs and actions which differs between the two treatments. In particular, the same history of play leads to different beliefs, and the same belief leads to different action choices, in finite and indefinite games. We then use the subjects' beliefs over actions in each round to identify their supergame beliefs over supergame strategies played by the other player. We find that these supergame beliefs properly capture the different classes of strategies used in each game. Importantly, subjects using different strategies have different supergame beliefs and for the most part strategies are subjectively rational given supergame beliefs. We also find that subjects underestimate the likelihood that others move to defection earlier than they do.

JEL classification: C72, C73, C92
Keywords: supergame, 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.

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, etc.—has been an active area of research that has produced many theories. However, our empirical understanding of the subject is much more limited, and specifically so in the controlled experimental setting: Although a recent experimental literature on the subject informs us of what behavioral patterns to expect in such an environment, the bulk of our knowledge concerns how the level of cooperation varies with such 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 in the decision making process to help us better understand behavior in such an environment.

A player's belief about other players' strategies forms the foundation of equilibrium analysis: it is assumed that beliefs correctly identify the strategies played by other players, and that the strategies best respond to those beliefs. In repeated games, however, strategies are complex since 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 thus face a practical challenge in forming a belief that correctly predicts other players' strategies. This is potentially made even more difficult by multiplicity of equilibria. When repeated, PD generates diverse patterns of dynamic behavior rationalized by different beliefs, and as a result 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 approach such repeated games. For instance, we may find evidence pointing to non-standard preferences if strategies don't best respond to beliefs, or to a failure of learning if the strategies do best respond to beliefs but beliefs are incorrect. We may also find evidence of beliefs that place positive probability on particular supergame strategies, which could explain the presence of cooperation in finitely repeated environments.

Comparison between finitely and infinitely repeated PD provides a useful instrument 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, experimental analysis has identified settings that induce 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 these two games.

There is a growing number of experiments with belief elicitation. (See the review of Schotter & Treviño (2014) and the references therein.) However, to the extent that such experiments have induced repeated games in the laboratory, they do not discuss the dynamic incentives of the players. The focus of the literature is either on the round-by-round evolution of beliefs over stage actions in games where the equilibrium payoff set does not expand with repetition, or on the conditional relationship of action choices with beliefs in each round where the repeated game is a byproduct of using fixed-pairing for subjects.[1] Our experiment is the first to perform belief elicitation in repeated games where dynamic incentives are clearly visible, and also study beliefs over supergame strategies rather than over stage actions.[2]

As a first foray into beliefs in repeated PD games, there are many questions that could be of interest. However, given the challenges associated with both implementing repeated games in the laboratory and eliciting beliefs, we have opted for simplicity whenever possible. Most importantly, we only elicit (first-order) beliefs about the other player's 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, where beliefs are front and center on the theoretical side.

Our experiment consists of two treatments referred to as *Finite* and *Indefinite games*. In the former, subjects play PD repeated over eight rounds; and, in the latter, subjects play PD over a random number of rounds with continuation probability of 7/8.[3] These parameters and the stage game were selected based on prior results in the literature: They are expected to generate not only significant levels of cooperation in both Finite and Indefinite games, but also similar levels of round one cooperation in both environments.[4] Our intention was to create two treatments where behavior was expected to be similar despite the theoretical difference. The treatment variation hence permits comparison of beliefs of subjects taking the same action in the same round (potentially along the same history) across Finite and Indefinite games, and provides insight into whether their strategic reasoning is similar or different across these two games.

[1] This literature is reviewed in more detail in Section A of the Appendix.

[2] One notable exception is a recent project by Gill and Rosokha that directly elicit beliefs over strategies.

[3] 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 associated with the discount factor of the infinitely repeated game.

[4] It is easy to find parameters that would generate very different initial cooperation rates between these two games (Dal Bó 2005).

We find that elicited beliefs accurately predict the other player's action along many histories. However, we also document some systematic deviations that can be tied to the slow unravelling of cooperation in Finite games. We also find that beliefs are *not* simply the (weighted) empirical average of past observations and that they are *forward looking*. For example, subjects correctly anticipate a decline in the likelihood that their opponent will cooperate in later rounds of Finite games.

The finding that beliefs are not simply the summary of past history implies that 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 that 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 that evolution of behavior can be well described by learning and evolutionary models over supergames (Dal Bó & Fréchette 2018, Embrey et al. 2017, Proto et al. 2020).

Our estimation approach assumes that each subject is endowed with a supergame strategy as well as a supergame belief.[5] 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.

The results clearly 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).[6] These observations suggest that heterogeneity in behavior can be explained, to a large extent, by heterogeneity in beliefs. More generally, in both Finite and 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 underestimate the likelihood that others move to defection earlier 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.

[5]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, then Bayesian updating will lead in the long run to the NE play of the repeated game.

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

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 $A = A_1 \times A_2$ 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 that player i believes that 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.*[7] One interpretation of $\tilde{p}_i$ is that it represents i's prior belief over the proportion of different strategies played by other subjects in that session.[8]

Note that $\tilde{p}_i$ 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

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

[7] We use $\tilde{p}$, instead of p to denote beliefs. In later sections, p is used to denote the actual distribution of strategies in the population.

[8] 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.

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 that 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 $\tilde{p}_i$:

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

Some of the key supergame strategies in our analysis are as follows. AC and AD are the strategies that choose C and D, respectively, for every history. σ_i is Grim if $\sigma_i^t(h^{t-1}) = C$ 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.[9]

3 Design

The experiment involves two (between-subjects) treatments which will be referred to as *Finite* and *Indefinite* games. Three important considerations (beside the aforementioned aim for simplicity) guided our experimental design.

1. *Observing high levels of initial cooperation in both Finite and Indefinite games (if possible at similar rates).* It is reasonable to expect people to have different beliefs in Finite and Indefinite games especially when these games generate very different behavior (cooperation rates). Although there can be value to documenting this, it is potentially more interesting to focus on the more puzzling case where these two types of games generate very similar behavior (specifically in terms of initial cooperation rates). 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

[9] All strategies considered in our analysis are listed and defined in Table 9 of the Online Appendix.

Table 1: Stage Game

In ECU

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

Normalized

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

two games. To achieve this we based our parameter selections on previous experiments and meta-analysis (Embrey et al. 2017, Dal Bó & Fréchette 2018).

2. *Introducing belief elicitation while mitigating the impact this 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.[10] The results of the current experiment reproduce the qualitative features of previous experiments without belief elicitation (with similar parameters). While beliefs at the start of the experiment are not elicited in the two-part approach, the potential benefits of not influencing behavior by discussing beliefs from the onset seemed to outweigh the downside of not observing beliefs in those early supergames.

3. *Allowing subjects to gain ample experience.* Prior research, both with finite and indefinite games show the importance of experience (Embrey et al. 2017, Dal Bó & Fréchette 2018).[11] For instance, for the parameters we use in Finite games, Embrey et al. (2017) find that the average round of last cooperation moves one round earlier for every ten supergames. This desire to have subjects play as many supergames as possible is one of the factors that support the need for simplicity. Asking more complex belief questions would slow down the experiment.

We now turn to the specifics of the experimental design.

The left panel of Table 1 shows the stage game used in the experiment (in experimental currency units) whereas the right panel shows its normalized version.[12] We use *supergame*

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

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

[12] The normalization facilitates comparison with prior studies. With normalization, we set mutual cooperation payoff equal to one and the mutual defection payoff equal to zero. The normalized temptation

to refer to each repeated game played between two matched players, and *round* to refer to each round of play within a supergame. In Finite games, each supergame ends after eight rounds, $T = 8$. In Indefinite games, after each round, there is a $\frac{7}{8}$ probability that the supergame will continue for an additional round.[13] In order 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.[14] In Indefinite games, *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 Indefinite games 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, while if a supergame lasts 10 rounds, $T = 10$.

At the conclusion of each supergame, subjects are randomly re-matched to play a new supergame. This process continues until the first supergame to terminate after at least one hour of play has elapsed. 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.[15] The task is incentivized via the *binarized scoring rule,* which determines the likelihood that a subject wins 50 experimental currency units based on their response in this task and the realized action choice of the matched subject.[16] The belief question is presented on a separate screen after subjects made action decisions for that round and before feedback is provided.

Two subjects with identical beliefs could make different choices because of different risk preferences. 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.[17]

We conducted eight sessions per treatment and sixteen sessions in total.[18] Table 2

payoff is hence $2 = (63 - 39)/(51 - 39)$ and the normalized sucker payoff is $-1.41 = (22 - 39)/(51 - 39)$.

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

[14]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.

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

[16]Unlike the classical quadratic scoring rule which 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 Berg et al. (1986) for an earlier formulation of the idea. We used the implementation outlined in Wilson & Vespa (2018).

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

[18]This is more sessions per treatment than typical. The reason for this will become apparent in the

Table 2: Session Summary

				No. of Game Rounds				Total no. of
		No. of	No. of	Actions	Actions and Beliefs			Obs.
Treatment	Session	Subjects	Supergames	Only	Early		Late	Rounds
Finite	1	20	12			8, 8,		96
	2	20	12			8, 8,		96
	3	20	13			8, 8, 8,		104
	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
Indefinite	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
	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).

summarizes basic information about each session. The supergames for the part with belief elicitation are separated into *early* and *late*. This will be used in the presentation of results, with most of the data analysis focusing on late supergames.[19] 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.[20]

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. Staring with Finite games (the left panel), we observe relatively high initial (round 1) 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. Cooperation rate $\bar{x}^8$ in round eight is decreasing with experience and low by the end (below 20%).

section on beliefs over strategies as the method we propose is data intensive.

[19]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.

[20]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 the Online Appendix C) were read aloud. The computer interface was implemented using zTree (Fischbacher (2007)) and subjects were recruited from UCSB students using the ORSEE software (Greiner (2015)).

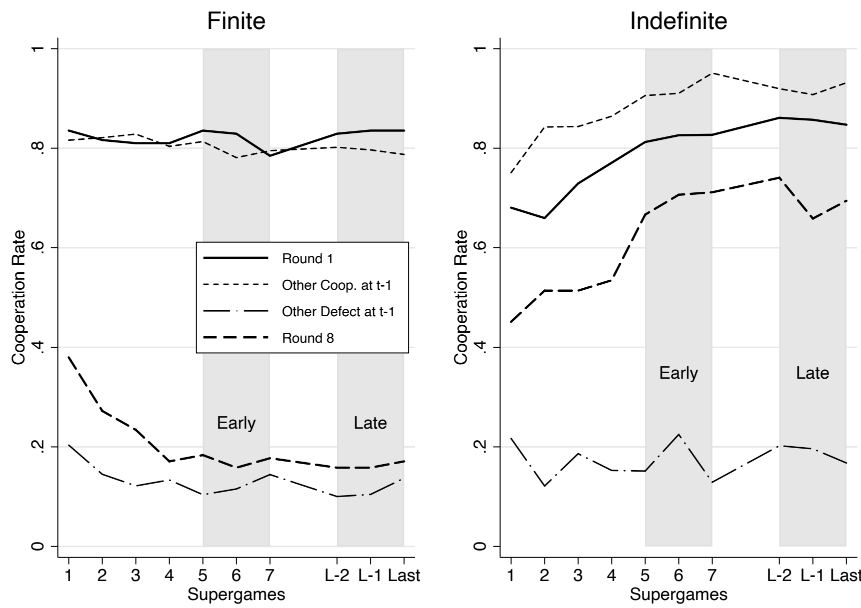

Figure 1: Cooperation Rate Over Supergames

The right panel of Figure 1 presents the same statistics for Indefinite games. In this case, and as with Finite games: Round one cooperation rates are high (start slightly below 80% and increase to slightly above 80%). Cooperation rates following cooperation by the other are high while cooperation rates following defection are low. Again, responsiveness increases with experience. However, in contrast to Finite games, cooperation rates in round eight are high and increasing with experience.[21]

Hence, consistent with prior experiments, the design successfully generates similar and high levels of round one cooperation in both games. Also in line with prior findings, subjects display responsiveness, and this increases with experience. Finally, cooperation collapses at the end of Finite games but persists in Indefinite games. In summary, behavior along key dimensions is qualitatively consistent with prior findings on these two games, and there is no indication of important changes in the subjects' behavior caused by the belief elicitation task in these environments.[22]

[21] Instead of round eight, one might want to compare the round eight behavior in Finite games to the last game round in Indefinite games, or to the last observation round. Doing so does not qualitatively change the results. These alternative figures are presented in the Online Appendix (Figure 13).

[22] Table 6 in the Online Appendix also shows that there are no significant changes in round one choices for supergames where beliefs are elicited (compared to those where they are not).

Result 1 We reproduce qualitative data patterns observed in previous experiments on Finite and Indefinite PD games. In particular, our results confirm that cooperation is history dependent in both games. Furthermore, cooperation evolves differently in both games: it collapses at the end only in Finite games.

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.

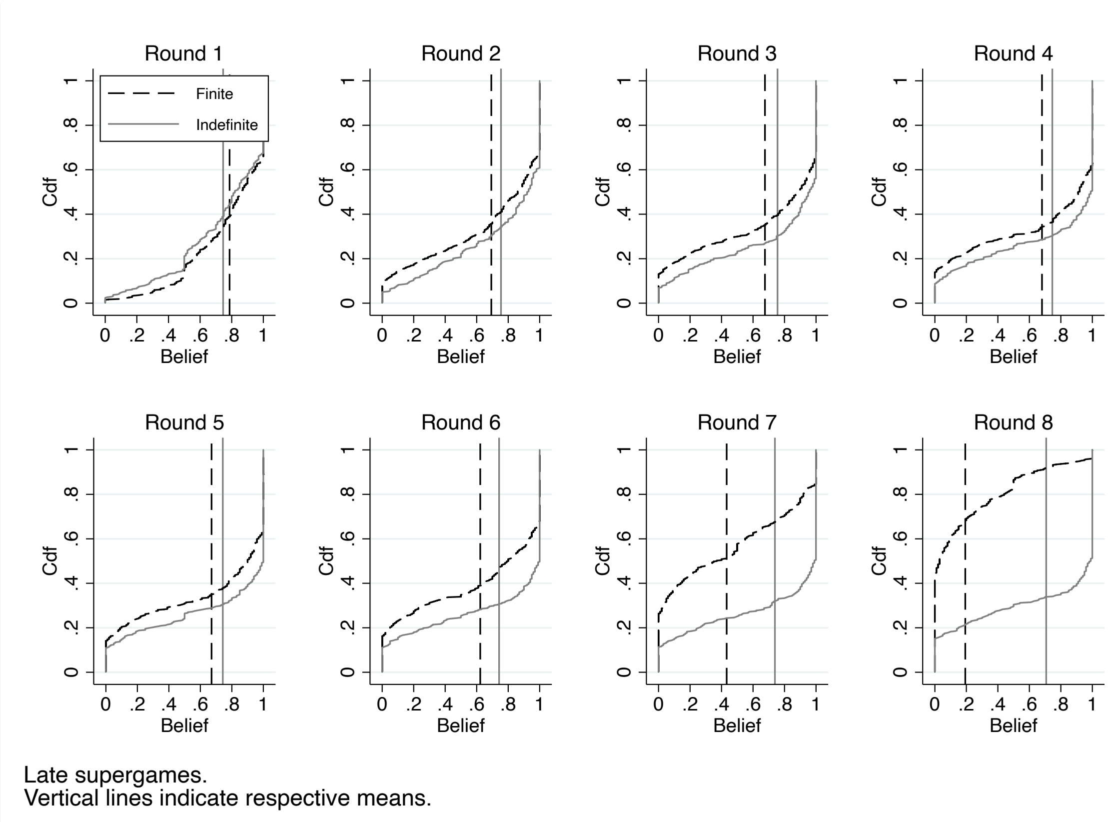

Figure 2: Distribution of Beliefs by Round

Figure 2 displays the cumulative distributions of (unconditional) beliefs in rounds $t = 1, \ldots, 8$ of late supergames in Finite and Indefinite games.[23] As the figure clearly shows

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

beliefs evolve very differently over rounds across these two types of games.[24] Beliefs become comparatively more pessimistic in Finite games as the supergame unfolds. The difference in average belief is statistically significant in rounds six, seven, and eight.[25]

A few more observations are worth making at this point. 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 Finite games, while more than 40% of subjects place probability one on cooperation by the other player in Indefinite games.

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

4.2.1 Actions and Beliefs

Putting beliefs and actions together reveals that 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.[26] Throughout the first eight rounds, the overall difference between action frequency and beliefs is less than one percentage point for Finite games and about two percentage points for Indefinite games. To the extent that they differ, the deviations are mostly for late rounds in Finite games and early rounds for Indefinite games. For example, the average round seven belief in Finite games $\bar{\mu}^7 = 43\%$ is more than nine percentage points higher than the observed cooperation rate of $\bar{x}^7 = 34\%$. In Indefinite games, we have the reverse pattern for round one, where the average belief $\bar{\mu}^1 = 74\%$ is 11 percentage points lower than the observed cooperation rate of $\bar{x}^1 = 85\%$.

When aggregated over all rounds, the overall difference between action frequencies and beliefs is not statistically significant for Finite games, but it is for Indefinite games (even though the difference is small in magnitude).[27] However, when we look at each round separately, both in Finite and Indefinite games, there is a statistical difference between

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

[25] Respectively $p < 0.05$, $p < 0.01$, and $p < 0.01$. Throughout, when (not) statistically significant is used without 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. (2017) for a discussion of issues related to hypothesis testing for experimental data). In the case of beliefs, as here, a tobit specification allowing for truncation is used. For tests of cooperation, a probit specification is used.

[26] As is well known, confidence intervals can overlap while two variables are statistically different, but not the opposite.

[27] This test is performed on the difference between the opponents action (coded as one for cooperate and zero for defect). Results are robust to including all observation rounds or only the first eight rounds.

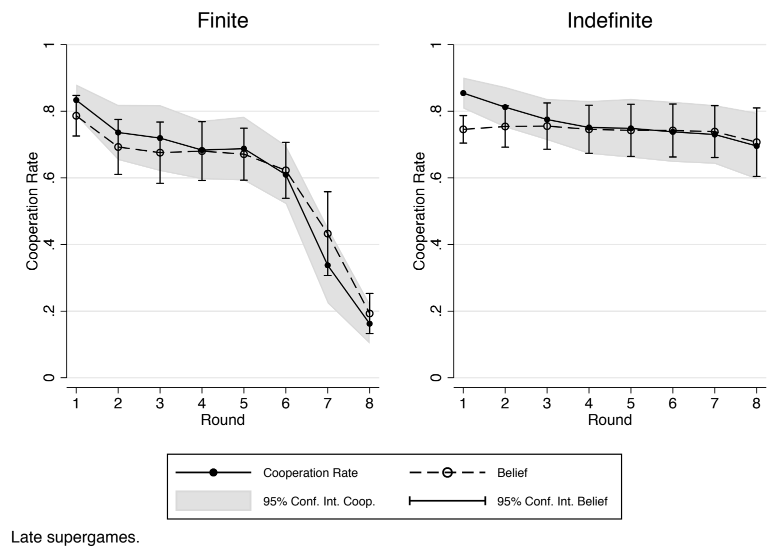

Figure 3: Choices and Beliefs by Round

action frequencies and beliefs for rounds one through three. The difference is about four percentage points for each of the three rounds of Finite games, while it is 11, 5.8 and 0.2 percentage points for the same rounds of Indefinite games. In rounds seven and eight there are also statistically significant differences between action frequencies and average beliefs for Finite games. In other rounds (rounds 4-6 of Finite games and rounds 4-8 of Indefinite games), beliefs and cooperation rates are not statistically different at the 10% level.

One natural question is whether, with experience, subjects learn to correct their mispredictions. Figure 4 displays the error in key rounds for early versus late supergames. As the figure shows, in many cases where there is a more substantial error in early supergames, there is improvement in late supergames: but not for round seven of Finite games and round one of Indefinite games. Even in these cases, however, subjects' beliefs do move in the right direction: As seen in Figure 15 in the Online Appendix 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 necessitate beliefs to be systematically off. For instance, in that same figure one can see that cooperation rates in round seven of Indefinite games are changing with experience but that this is correctly anticipated by subjects as reflected in their beliefs.

Although determining exactly how beliefs are formed is not the goal of this study,

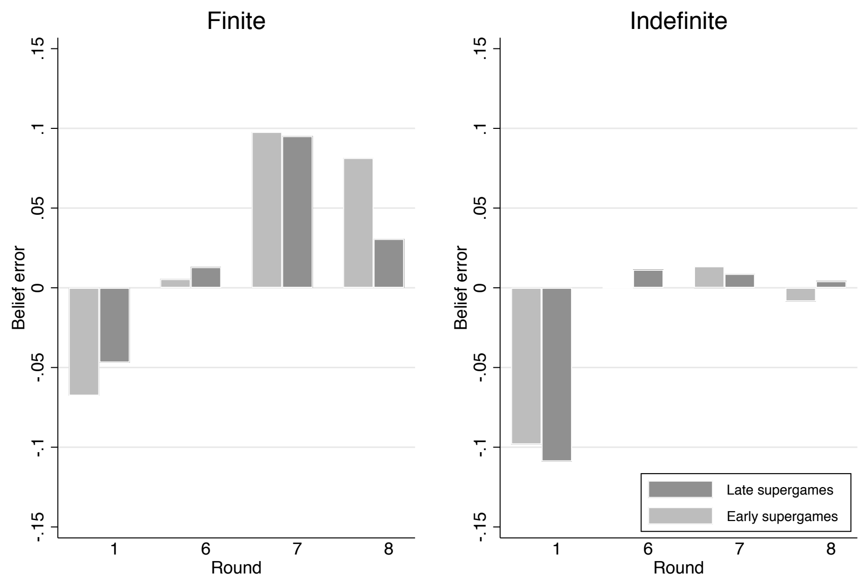

Belief error denotes average difference between beliefs and actions.

Figure 4: Belief Errors in Early vs. Late Supergames

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 introspection alone, or some combination of learning and introspection.[28] The data suggests that although experiences matter in shaping beliefs, they are not the sole determinant. Figure 16 in the Online Appendix 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 only unconditional beliefs were considered, but what about 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 presents the relevant data conditional on round one histories (labeled with one's own action first followed by the opponent's action). In both Finite and Indefinite games, we

[28]The earlier observation about Indefinite games— although behavior is changing in round seven, beliefs track action frequencies closely—already suggests that subjects cannot be basing their beliefs only on empirical frequencies.

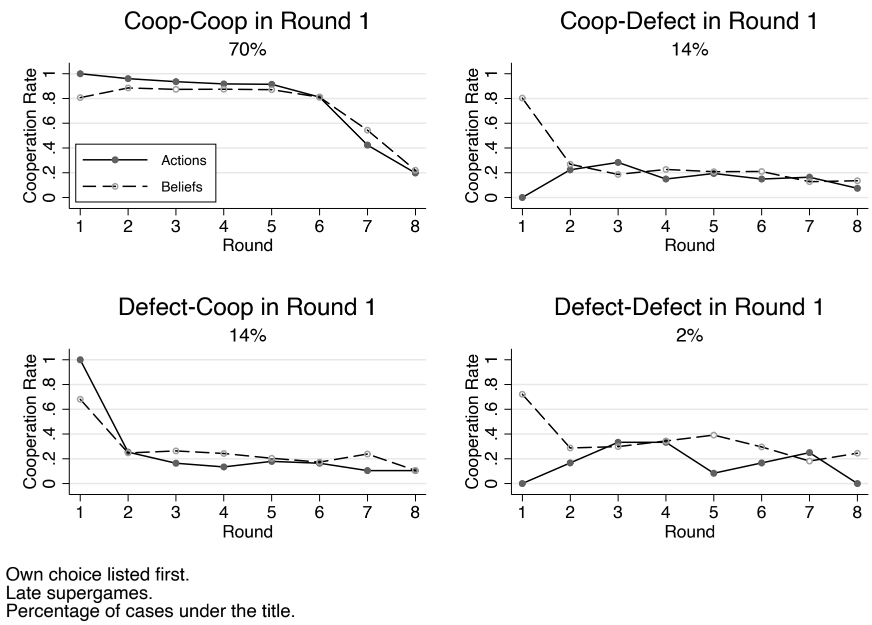

Figure 5: Conditional Round 2 Beliefs, Finite Games

observe that beliefs quickly adjust in response to the other's action.[29] It is interesting to note that the downward adjustments following a unilateral choice of D by either player in round 1 are of the correct magnitude even though the required adjustment is quite large. Comparing across the two figures, 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 Finite games, most of the initially cooperative interactions eventually break down, and this is mirrored by beliefs. In Indefinite games, on the other hand, beliefs on cooperation are sustained if they survive the second round.

These results about beliefs being 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,

[29]They should not be correct in round one since beliefs are unconditional while by construction the figures present specific action frequencies in round one.

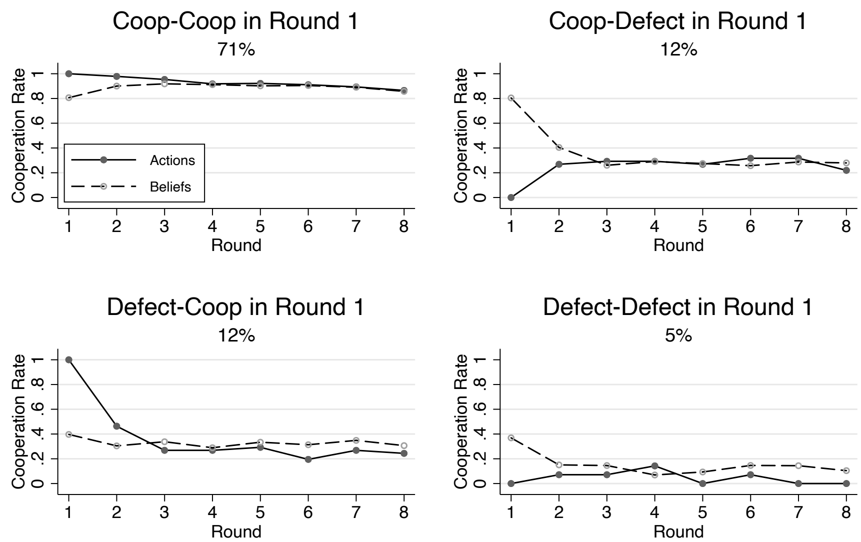

Own choice listed first.
Late supergames.
Percentage of cases under the title.

Figure 6: Conditional Round 2 Beliefs, Indefinite Games

and *uncertain* if the empirical frequency is in 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. This is done for rounds one and two.

Table 7 in the Online Appendix shows that accuracy of beliefs at the individual level, as defined above, is high both for round one (73% in Finite games, 67% in Indefinite games) and round two (83% in Finite games, 80% in Indefinite games). The accuracy rate is substantially above 33% (the benchmark if beliefs were random) even in early supergames. However, there is one history after which accuracy is low: In round two of Indefinite games 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 7 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 such as $h^1 = (C, C)$), but less so along other histories that are less common (particularly along $h^1 = (C, D)$ and (D, C) in Indefinite games).

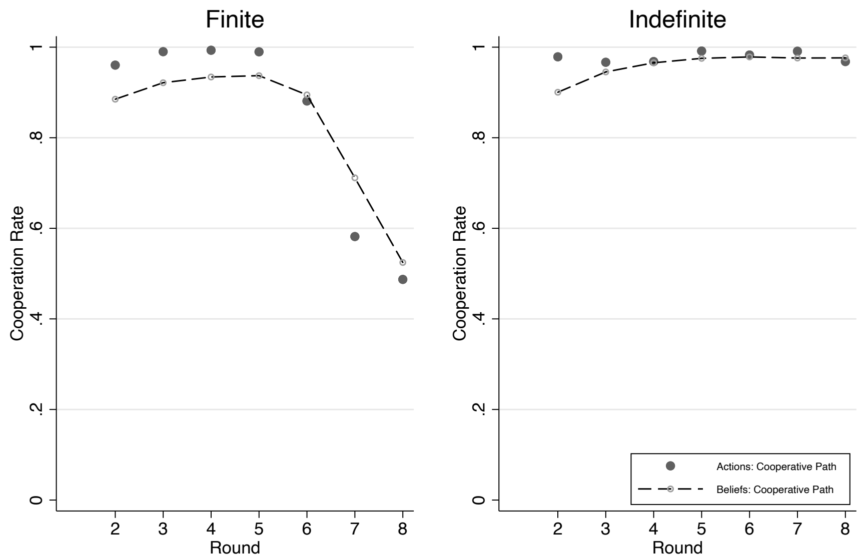

Cooperative path: joint cooperation up to that round.
Late supergames.

Figure 7: Cooperative Path (First Eight Rounds)

As the Figures 3 and 4 above show, supergames starting with joint cooperation are the most common. How do beliefs evolve on a mutual cooperation path? Figure 7 shows the average cooperation rates $\bar{x}^t$ and average beliefs $\bar{\mu}^t$ along the history $h^{t-1} = ((C,C),\ldots,(C,C))$.[30] 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 towards the end in Finite games. 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.[31] Nonetheless, we see clear evidence that subjects underestimate the degree to which cooperation drops from round 6 to 7: while 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.[32] In summary, these findings suggest that while subjects anticipate the decline in cooperation, they underestimate the magnitude and foresee only 60% of the actual drop in cooperation. In Indefinite games, on the other

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

[31]This 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.

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

hand, beliefs and cooperation rates remain high as the supergames unfolds. We note also that these patterns are already visible in early supergames (see Figure 17 in the Online Appendix).

The last observation suggests in particular that the evolution of beliefs in Finite games cannot simply be explained by heuristic models based on past action choices (within a supergame). For example, if a subject always set his belief equal to his opponent's action in the previous round, then he would report beliefs for round 7 (in Finite games) that are almost 3 times more over-optimistic and less than half as accurate than the ones we observe in the data.[33] Clearly, beliefs in Finite games 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 Finite games and pessimistic early in Indefinite games. (2) Beliefs respond to history of play. (3) There are, however, differences across games even when conditioned on the same history. In particular, subjects correctly anticipate that cooperation will breakdown despite a history of joint cooperation in Finite games.

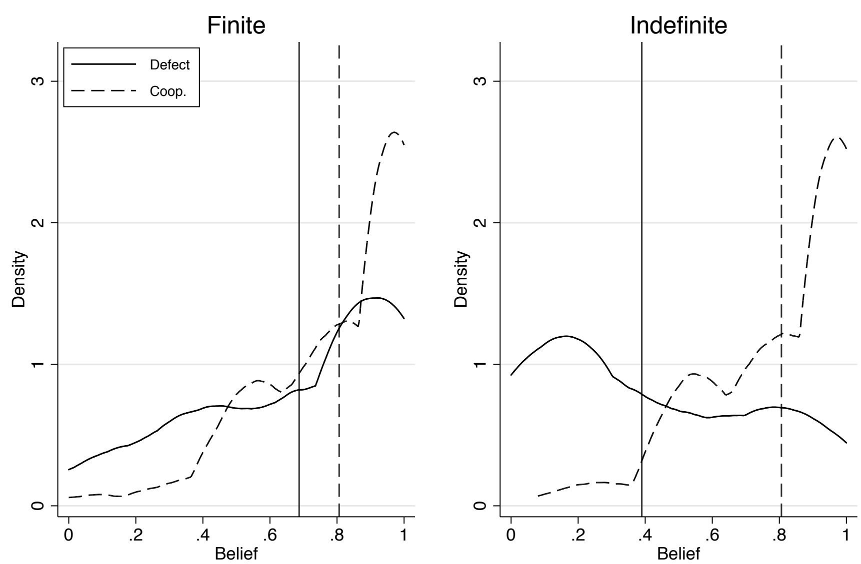

Vertical lines indicate respective averages.
Late supergames.

Figure 8: Beliefs of Defectors vs Cooperators in Round 1

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

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 1 beliefs μ_i^1 by treatment and by the subject's own action a_i^1 in round 1. At a broad level, it is easy to see that beliefs of cooperators and defectors are more different from one another in Indefinite games than in Finite games. The average beliefs of cooperators and defectors are statistically different in Indefinite games ($p < 0.01$) but not in Finite games.[34] Of those subjects who reported a belief less than 50% in round 1, only 39% cooperated in Indefinite games in contrast with 55% who cooperated in Finite games. Subjects with optimistic beliefs cooperated in both treatments: of those subjects who reported a belief greater than 50% in round 1, 94% cooperated in Indefinite games and 87% cooperated in Finite games. In other words, round 1 beliefs were more predictive of actions in Indefinite games than in Finite games and subjects with higher beliefs tend to defect more often in Finite games than in Indefinite games.

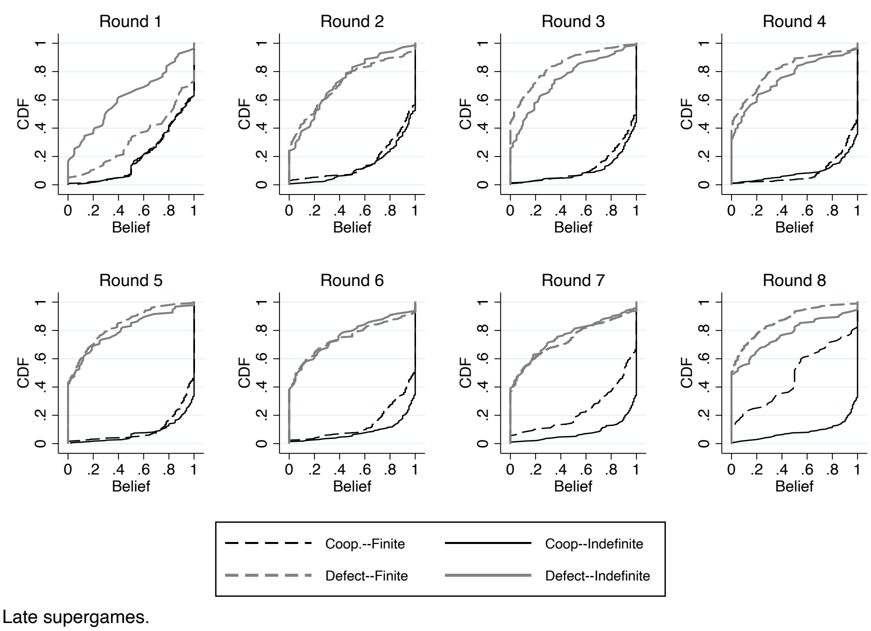

Figure 9: Beliefs by Action and Treatment: Rounds 1-8

Figure 9 plots the CDF of beliefs by action and treatment for each round. It clearly shows that cooperation and defection are associated with different beliefs. Except for round one in Finite games, in every other comparison—every round for each treatment—

[34]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.

the average belief is statistically different between those who cooperate versus those who defect ($p < 0.01$). Higher cooperation rates are associated with more optimistic beliefs more generally. Table 8 in the Online Appendix shows that in all rounds, marginal impacts of beliefs on the likelihood of cooperation are positive in both Finite and Indefinite games. In addition, the round number has a significant negative impact on cooperation in Finite games but is insignificant in Indefinite games.[35] More specifically, subjects are more likely to defect later in Finite games, even if their beliefs are the same.

Perhaps more surprisingly, cooperation and defection in key 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 ($p < 0.01$) and so are those of the subjects who defect ($p < 0.1$). Subjects who defect in round eight of Finite games are more pessimistic (on average) than those who do so in Indefinite games. Similarly subjects who cooperate are more optimistic in Indefinite games than those in Finite games. On the other hand, subjects who defect in round one of Finite games are more optimistic than those who do so in Indefinite games ($p < 0.01$). Hence the same actions 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 that there exists a link between beliefs and actions, and also that beliefs are not just the summary of past action choices in a supergame. This leads us to the consideration of beliefs over strategies.[36] 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 how strategies best respond to beliefs.

Plan for 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.

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

[35] Specifications with our measure of risk attitude finds that those are not statistically significant.

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

over strategies for one player (referred to as player 1) when the data available to us is his round beliefs over actions elicited in one supergame (against player 2). For the purpose of the example, assume that we know player 1 considers player 2 to be using 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, we can already infer the probability player 1 associates with player 2 playing AD, since that is the only strategy 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. For that we look at beliefs elicited in other rounds of the supergame. Assume in round one player 1 plays D and player 2 plays C. After observing this history, player 1 reports his round two belief: $\mu_1^2 = 0.1$. Since player 2 started by playing C, player 1 now knows she is not playing AD. However, player 1's belief on 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 of how we can recover ex-ante beliefs over strategies using data on beliefs over stage actions elicited in each round of a supergame. To this we add that players 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 in adopting it to our dataset. 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 there is a significant body of literature documenting consistent patterns on 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 could 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 that 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 that subjects of the same type share the same beliefs.[37]

[37] To explore whether subjects of the same type have similar beliefs, we do the following exercise. We compute 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 types relative to all others in the population. Out of the 12 types observed in Finite games and the 11

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 strategies used in the population of subjects.[38] 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 rely on prior evidence to select our set of candidate strategies. Specifically, we include the nine strategies that are statistically significant according to in the SFEM estimates of Fudenberg et al. (2012) in any of their nine Indefinite PD games with perfect monitoring. We also include all possible threshold strategies: The threshold strategies are identified as the key strategies that emerge with experience in Finite PD games (Embrey et al. 2017). This results in the inclusion of a total of 16 strategies: AD, AC, Grim, Grim2, Grim3, TFT, TF2T, 2TFT, STFT, T2–T8.The estimation is performed on late supergames.

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 Finite games (Embrey et al. 2017), and AD, Grim, and TFT usually account for a majority of the strategies in Indefinite games (Dal Bó & Fréchette 2018).

More specifically, we note that, in Finite games, 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%. Beside these, TFT and Grim, are the other commonly used strategies (at the 4th and 6th positions). When played against other strategies in the top 5, these strategies will cooperate in round one, and then depending on the history of play, either cooperate all the way or defect in the last few rounds.

In Indefinite games, conditionally cooperative strategies dominate, with TFT and Grim representing almost half of the choices. This is followed by the two unconditional strategies, AC and AD. Together these four account for more than two thirds of the strategies. The lenient versions of Grim and TFT are the next most popular, together 22% of the choices. All other strategies are at most 5% each, and the threshold strategies are almost completely irrelevant. Together, conditionally cooperative strategies account for 80% of the data (by

types in Indefinite games; only one of the 23 paired comparisons is not in line with the assumption that the spread in beliefs is less among subjects of the same type.

[38]See Dal Bó & Fréchette (2018) for results across a variety of experiments applying the SFEM. Dal Bó & Fréchette (2019) show that the SFEM produces results comparable to direct elicitation.

Table 3: Strategy Prevalence and Typing

Finite			Indefinite		
	Share			Share	
Type	SFEM	Typing	Type	SFEM	Typing
T7	0.30	0.35	TFT	0.34	0.58
T8	0.22	0.20	Grim	0.15	0.07
AD	0.12	0.12	AC	0.10	0.10
TFT	0.09	0.12	AD	0.09	0.10
T6	0.08	0.08	TF2T	0.09	0.03
Grim	0.07	0.02	Grim2	0.07	0.02
TF2T	0.03	0.04	Grim3	0.06	0.02
Grim3	0.03	0.03	2TFT	0.05	0.01
STFT	0.02	0.02	STFT	0.04	0.04
AC	0.02	0.01	T3	0.02	0.03
Grim2	0.01	0.01	T8	0.01	0.01
T2	0.01	0.01	T7	0.00	0.00
2TFT	0.00	0.00	T6	0.00	0.00
T5	0.00	0.00	T5	0.00	0.00
T4	0.00	0.00	T4	0.00	0.00
T3	0.00	0.00	T2	0.00	0.00

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

contrast, these represent only 25% of the Finite data).

Result 5 We reproduce results about strategic choices observed in previous Finite and Indefinite PD games.[39] In particular our results confirm that there is strategic heterogeneity within and across treatments. In Finite games, subjects mostly use threshold strategies; while in Indefinite games, 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 that SFEM yields the probability distribution over supergame strategies (p) and the probability of implementation errors ($1 - \beta$). These 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

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

that has the highest likelihood according to this posterior.[40]

To demonstrate how this works, consider a simpler setup where the set Z of candidate strategies consists only of AD and AC. Assume that 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 that the strategy of each subject is chosen from the set $\widehat{Z} = \{\widehat{AD}, \widehat{AC}\}$ using the prior distribution p.[41] Assume now that there is a subject who, over multiple supergames consisting of 24 rounds in total, cooperates in 20 rounds and defects in 4 rounds. Given p and β, we can calculate the Bayesian posterior that this subject is playing $\widehat{AD}$ vs. $\widehat{AC}$. 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, while 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 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 Indefinite games, the fraction of subjects typed as TFT is greater than the fraction of TFT in the population.[42] Clearly, the smaller the fraction of subjects of a given type, the less reliable their belief estimates will be.

[40] There was a unique strategy within the consideration set for each subject in our data set that achieved the highest posterior (given the SFEM results).

[41] 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.

[42] Here are two potential sources of differences. First, and simply mechanically, some subjects play more supergames than others; and thus the fraction of subjects corresponding of a type can differ from the population (over supergames) fraction of that strategy. Second, imagine a data set where a large fraction of subjects play TFT, while a small fraction plays Grim. However, for some of the subjects playing Grim, the number of observations that distinguishes Grim from TFT is very few. 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.

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.[43] 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 that $\tilde{p}$ and $\tilde{\beta}$ jointly pin down beliefs over stage actions given each history. For illustration, suppose again that the set Z of candidate strategies consists only of AD and AC so that $\tilde{Z}^{\tilde{\beta}}$ consists of their randomized versions $\widetilde{\text{AD}}$ and $\widetilde{\text{AC}}$ for $\tilde{\beta} = 0.9$. It then follows that the round 1 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

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

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 that 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 is reported in Tables 4 and 5 with the complete results provided in the Online Appendix. Note that some types are not observed frequently enough to allow for estimation. This is the case whenever only one percent of subjects are of a certain type. In addition, especially for strategies that few subjects were typed as using, there is sometimes not enough variation to separate the beliefs on some of the strategies. In those cases, we set the least popular strategies (according to SFEM) to zero, and "assign" the belief to the more popular strategy. This applies to only three of the 84 estimates reported in Tables 4 and 5.

It is immediately clear from these belief estimates that supergame beliefs are substantially heterogeneous both across different types and across games. Interestingly, given that these estimates are from late supergames, an implication is that experience does not eliminate this heterogeneity.

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

Table 4: Beliefs Over Strategies in Finite Games

	Share		Estimated Beliefs - $\tilde{p}$									
Type	SFEM	Typing	T8	T7	TFT	2TFT	GRIM	AD	TF2T	Other	ν	$\tilde{\beta}$
T7	0.30	0.35	0.43	0.43	0.00	0.00	0.14	0.00	0.00	0.00	0.04	1.00
T8	0.22	0.20	0.51	0.00	0.01	0.01	0.00	0.07	0.17	0.22	0.04	1.00
AD	0.12	0.12	*[0.00]*	0.00	0.00	0.76	*[0.00]*	0.23	0.00	0.00	0.06	1.00
TFT	0.09	0.12	0.30	0.00	0.55	0.00	0.00	0.11	0.00	0.04	0.05	1.00
T6	0.08	0.08	0.48	0.50	0.00	0.00	*[0.00]*	0.01	0.00	0.01	0.03	1.00
GRIM	0.07	0.02	0.23	0.00	0.26	0.00	0.00	0.17	0.05	0.28	0.07	1.00
Other	0.12	0.11	0.04	0.18	0.36	0.00	0.24	0.04	0.02	0.13		
All	0.00	0.00	0.33	0.19	0.12	0.09	0.07	0.07	0.04	0.09		

Estimation on late supergames out of 16 strategies: AD, AC, Grim, Grim2, Grim3, TFT, TF2T, 2TFT, STFT, T2–T8.
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 10.

Table 5: Beliefs Over Strategies in Indefinite Games

	Share		Estimated Beliefs - $\tilde{p}$									
Type	SFEM	Typing	Grim	TF2T	TFT	AC	AD	Grim2	STFT	Other	ν	$\tilde{\beta}$
TFT	0.34	0.58	0.21	0.14	0.27	0.08	0.07	0.16	0.04	0.02	0.01	1.00
Grim	0.15	0.07	0.74	0.08	0.03	0.09	0.00	0.06	0.00	0.00	0.06	1.00
AC	0.10	0.10	0.00	0.48	0.00	0.52	0.00	0.00	0.00	0.00	0.05	1.00
AD	0.09	0.10	0.07	0.00	0.01	0.01	0.90	0.01	0.00	0.01	0.04	1.00
TF2T	0.09	0.03	0.09	0.91	0.00	0.00	0.00	0.00	0.00	0.00	0.01	1.00
Grim2	0.07	0.02	0.00	0.24	0.00	0.24	0.00	0.52	0.00	0.00	0.05	1.00
Other	0.16	0.10	0.15	0.09	0.25	0.09	0.14	0.11	0.02	0.17		
All			0.22	0.22	0.14	0.13	0.13	0.12	0.02	0.03		

Estimation on late supergames out of 16 strategies: AD, AC, Grim, Grim2, Grim3, TFT, TF2T, 2TFT, STFT, T2–T8.
Rows, top 6 *played* strategies. Columns, top 7 *believed* strategies.
SFEM estimate for β is 0.94. Complete results in Table 11.

The two most common types in Finite games, T7 and T8, correctly anticipate that threshold strategies are the most common strategies. In fact, the average belief for all subjects correctly anticipates that these two strategies together represent 52% of the data. However, both T7 and T8 types display a systematic bias, namely they expect others to be less likely to defect before them relative to the truth. For instance, T7 overestimates the prevalence of T7, who would defect at the same time as themselves, as well as T8 and

Grim who would defect after them. In fact, they expect no one to defect before them by placing zero weight on AD or T6. Similarly, T8 overestimates the prevalence of others like them who will defect in round 8 as well as conditional cooperators, while underestimating the likelihood of others who would defect from the start. T6 believes that almost everyone plays T7 or T8. This generalized pattern can be observed in the left panel of Figure 10, which displays the belief that others will defect before, after, or at the same time as oneself compared with the reality.[44]

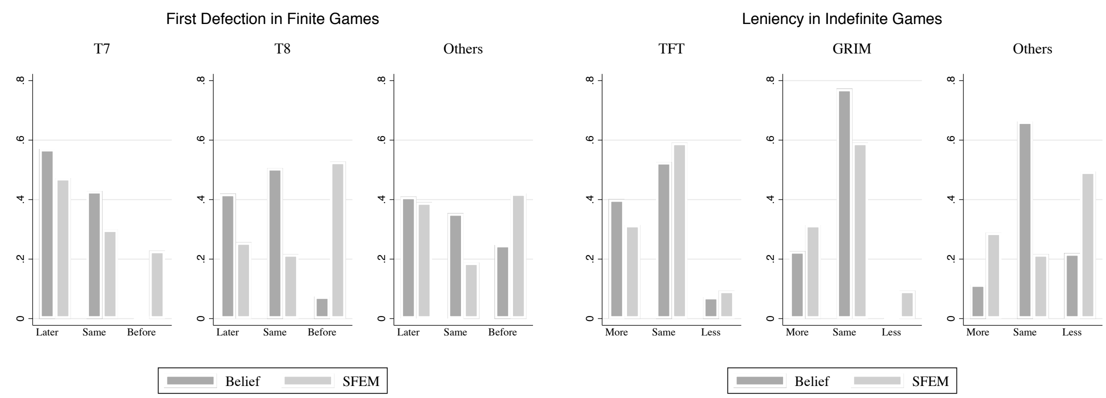

Figure 10: Beliefs vs. Reality

In Indefinite games, subjects overall correctly anticipate that others mostly adopt conditionally cooperative strategies, with their beliefs over those strategies (on average) totaling 72%, which is close to their total SFEM share at 80%. The most common type, TFT, has a quite accurate assessment of the total popularity of the most common strategies, TFT and Grim, estimating their total to be 48%, close to the actual 49%. However, they overestimate the fraction of Grim and underestimate the proportion of TFT. The second most popular type, Grim, substantially overestimates the fraction of Grim, while underestimating the popularity of TFT and a class of strategies that start by defection. The right panel of Figure 10 presents a similar—although less stark—picture to the right panel, focusing on leniency.[45] Relative to SFEM shares, subjects in Indefinite games overestimate the degree to which others are equally lenient or more lenient than they are. Hence, in both treatments, subjects display a tendency to believe others would not defect more, or first, or “as quickly”, than they do when compared to the actual distribution.

Note that ν provides a measure of how far reported beliefs are from those beliefs implied

[44]We adopt the following first-defection ranking over strategies in our consideration set: $AD \sim STFT \succ T2 \succ T3 \succ T4 \succ T5 \succ T6 \succ T7 \succ T8 \sim AC \sim Grim \sim Grim2 \sim Grim3 \sim TFT \sim TF2T \sim 2TFT$.

[45]We adopt the following leniency ranking over strategies in our consideration set: $AC \succ Grim3 \succ Grim2 \sim TF2T \succ Grim \sim TFT \sim STFT \sim 2TFT \sim T2 - T8 \succ AD$.

by $\tilde{p}$, $\tilde{\beta}$ and Bayesian updating. The estimated values for ν are quite low for the most common types: $\nu = 0.04$ for T7 in Finite games and $\nu = 0.01$ for TFT in Indefinite games. To interpret these values, consider $\nu = 0.05$. If the belief implied by the model $p^{\star}$ is 0.5, according to our functional form assumptions, the probability that the subject reports a belief $p \in [0.4, 0.6]$ is 0.76, the probability that $p \in [0.3, 0.7]$ is 0.96.[46] This implies that reported beliefs are not too far from those estimated by the model. Overall, this suggests that the model captures the variation in the reported beliefs quite well.

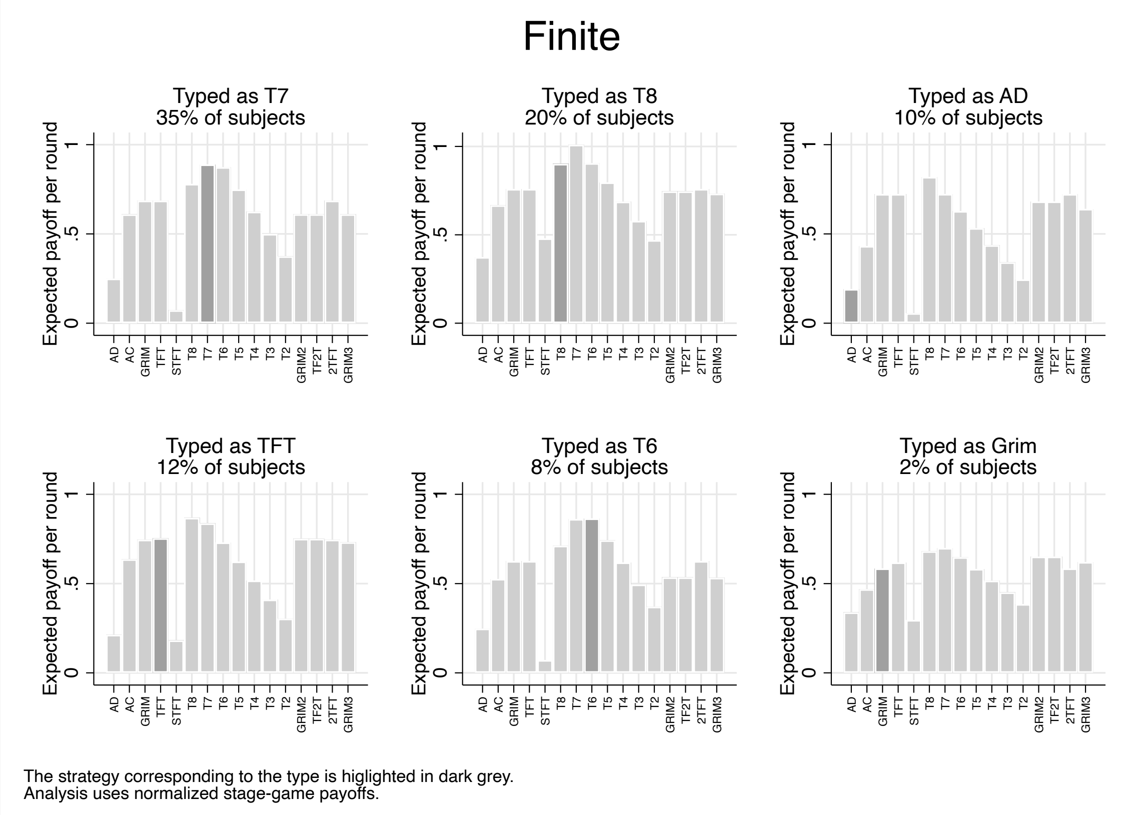

Figure 11: Best Response for Top 6 Types in Finite Games

In summary, two patterns emerge so far: 1) Overall, subjects' beliefs correctly capture important underlying differences between the two treatments: Their beliefs place positive weight on threshold strategies in Finite games and conditionally cooperative strategies in Indefinite games. 2) They tend to believe that others are less likely to defect before them than implied by the actual distribution.

[46]The corresponding values are 0.93 and almost 1 if $\nu = 0.03$.

A natural next question is whether the types thus identified are subjectively rational: Do their supergame strategies best respond to their supergame beliefs? We have so far imposed no restriction on the link between the strategies and beliefs, since 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 purpose of our discussion in this section, we say that a type is subjectively rational (given the assumed preferences) if its strategy choice is a best response to its supergame belief among those strategies Z in the consideration set.[47]

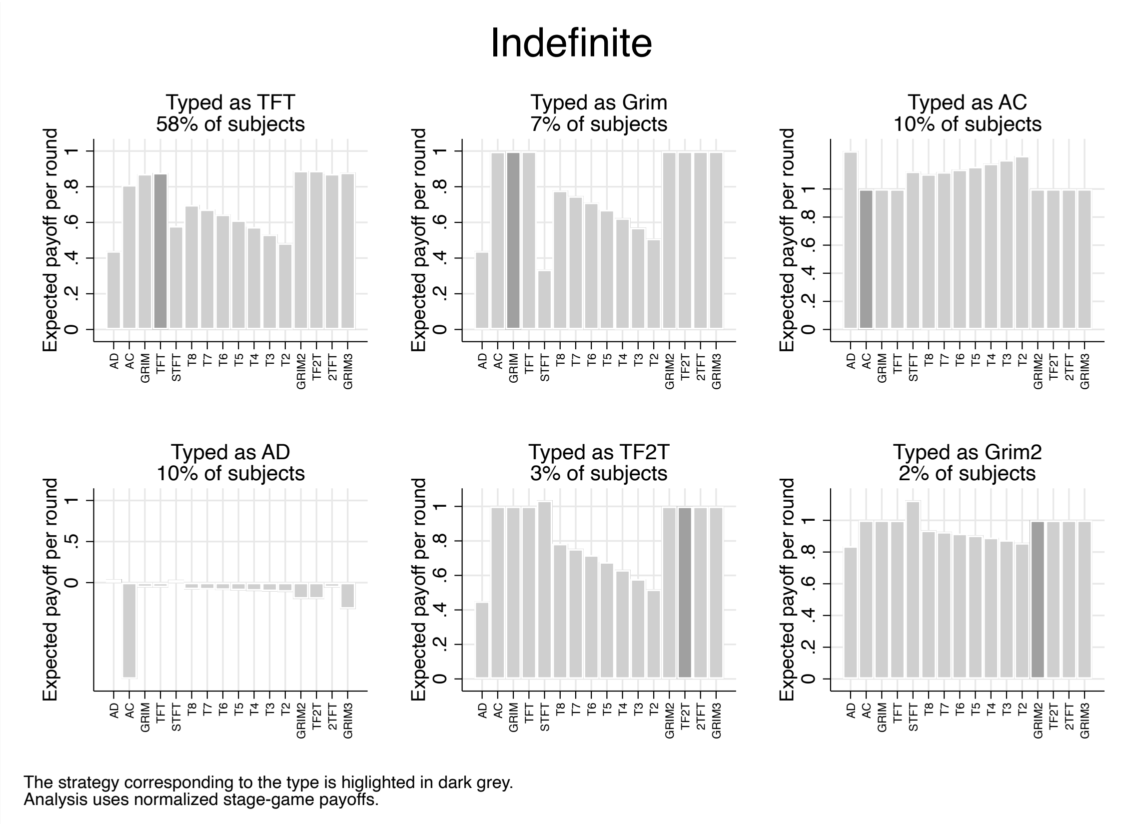

Figure 12: Best Response for Top 6 Types in Indefinite Games

The results, presented in Figures 11 and 12 suggest that most subjects' strategy choices

[47]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, since estimated values for $\tilde{\beta}$ are very close to 1, this has very little bearing on the results. To calculate the expected payoff of each strategy, we simulate play in 1000 supergames given $\tilde{\beta}$.

are either exact or approximate best responses given their supergame beliefs.[48] In Finite games, T7 and T6 types (43% of subjects) exactly best respond to their supergame beliefs, and T8, TFT, and Grim types (34% of subjects) approximately best respond to their supergame beliefs by obtaining 88% of the best response payoff on average. The only type whose strategy is far from a best response is AD (10%). In fact, their strategy choice is close to being the worst given the stated beliefs.[49]

In Indefinite games, a similar pattern emerges. Three of the most common types (TFT, Grim, and AD—75% of subjects) almost exactly best respond to their beliefs.[50] Likewise, TF2T and Grim2 (5% of subjects) approximately best respond to their expected payoffs. One "major" type far from best responding to their belief is AC (10%), 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. Hence, overall, the majority of subjects appear subjectively rational or close to it.

Note that 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 reveals T6 to be the best response to the population in Finite games and Grim2 to be the best response to the population in Indefinite games. In Finite games, the most frequent T7 type achieves 97% of the best response payoff from T6. In Indefinite games, we see that the most frequent TFT type achieves 93% of the best response payoff from Grim2. However, there are strategy-types that are much farther away from best responding to the population. For example, the AD type in Finite games only achieves 60% of the best response payoff.

Result 6 Subjects using different strategies have different supergame beliefs. Most types are close to best responding to their beliefs: they are subjectively rational. Subjects correctly anticipate the type of strategies that are popular in each treatment. In both games, however, subjects underestimate the likelihood that others move to defection earlier than they do.

[48]Table 12 in the Online Appendix provides detailed best response analysis for each of the six common types in both Finite and Indefinite games.

[49]Note that subjects playing AD receive weakly higher payoffs any supergame then their opponent. These subjects have little chance to observe what would happen along alternative histories. Hence, based on their experience, they might mistakenly believe they are making good choices. Proto et al. (2019) and Proto et al. (2020) provide related evidence on the importance of direct experience for learning the value of cooperation.

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

5 Conclusion

Beliefs play a central role in equilibrium theory, and there is increasing evidence that they are also key to the understanding of behavior observed empirically in repeated settings. This study measures 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 interact in this important class of games.

Our first key finding is that beliefs are, in aggregate, remarkably accurate. In both Finite and Indefinite games, beliefs averaged over all rounds are less than three percentage points away from the empirical action frequencies. Furthermore, the qualitative patterns of behavior—the persistence of cooperation in Indefinite games and its collapse toward the end in Finite games—are well anticipated by beliefs. Beliefs also adjust appropriately to the history of play even when these adjustments are not small: Along some round one histories, they move by almost 60 percentage points. While beliefs are heterogeneous with a varying degree of accuracies, a vast majority of subjects make correct predictions when judged based on a ternary partition over the likelihood of cooperation by the other player (likely, unlikely, or uncertain) conditional on the history of play. Moreover, beliefs are more accurate after the most common histories. To the extent that average beliefs do not correspond to the empirical action frequencies, the key deviations are over-optimism in late rounds of Finite games and over-pessimism in early rounds of Indefinite games.

Importantly, results show that beliefs over stage actions are *forward looking*. Most notably, beliefs along the history of mutual cooperation evolve very differently in Finite and Indefinite games. Persistence of cooperation in Indefinite games and its collapse in Finite games 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 finding is based on the development of a novel method to recover beliefs over supergame strategies from beliefs over stage actions in each round. This allows us to study the connection between the choice of supergame strategies and beliefs over them. In both Finite and Indefinite games, we observe that subjects playing different strategies have strikingly heterogeneous beliefs over the strategy choice of the other player. The results also show that subjects tend to underestimate the likelihood that others will move to defection before they do, and that for many strategies, their choice is subjectively rational in the sense that they generate an expected payoff close to the best response payoff against their beliefs. This in particular suggests that the observed heterogeneity in strategy choice is in part explained by the heterogeneity in beliefs. Among other things, the identified patterns provide insights into the forces that slow down unravelling of cooperation in Finite games.

In summary, our results on beliefs suggest that 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 that 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.

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.

Berg, J. E., Daley, L. A., Dickhaut, J. W. & O'Brien, J. R. (1986), 'Controlling Preferences for Lotteries on Units of Experimental Exchange', *The Quarterly Journal of Economics* **101**(2), 281–306.

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.

Embrey, M., Fréchette, G. R. & Yuksel, S. (2017), '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.

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.

Schotter, A. & Treviño, I. (2014), 'Belief elicitation in the laboratory', *Annual Review of Economics* **6**, 103–128.

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 NE played in fixed and random pairing, Nyarko & Schotter (2002) find that 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.[51] Hyndman et al. (2010) study beliefs when a stage game with a unique mixed Nash equilibrium is repeated 20 times, and find that 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*) discuss that some players attempt to influence the beliefs of other players through their own actions and as a result help the process converge to a NE.[52]

The experimental literature on beliefs examines the question of whether actions are best responses to beliefs with no definite answers. Nyarko & Schotter (2002) find that the actions in each round mostly best respond to the stated beliefs, but also find that fictitious play beliefs better predict the opponents' action than the stated beliefs. Costa-Gomes & Weizsäcker (2008) use fourteen 3×3 games to investigate the relationship between subjects' elicited beliefs and their strategy choice. Regardless of whether belief elicitation precedes strategy choice or not, Costa-Gomes & Weizsäcker (2008) find that 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 that 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 that 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 that 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).[53]

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

[52] 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.

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

Neugebauer et al. (2009) confirm this relationship in their experiment on a finitely repeated 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. (2017) find some key patterns by re-analyzing data from a collection of laboratory experiments.[54] First, in finitely repeated PD, the fraction of threshold strategies increases with experience.[55] 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.[56] In finite games, if the parameters are conducive to cooperation, round one cooperation increases with experience whereas last round cooperation decreases with it.[57] 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).[58] AD, Grim, and TFT are generally the most popular among them, 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.[59] 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 while it is insignificant in 20 paper-treatments.[60]

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

[55] 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.

[56] Embrey et al. (2017) 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.

[57] 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.

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

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

[60] This analysis eliminates all data in within-subjects designs after a change in treatment and only pre-

serves the initial treatment. Most of the insignificant cases have a small number of observations. There is one treatment with a negatively significant impact. This may be in part because of relatively low round one cooperation at 0.36.

B Additional Figures and Analysis

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

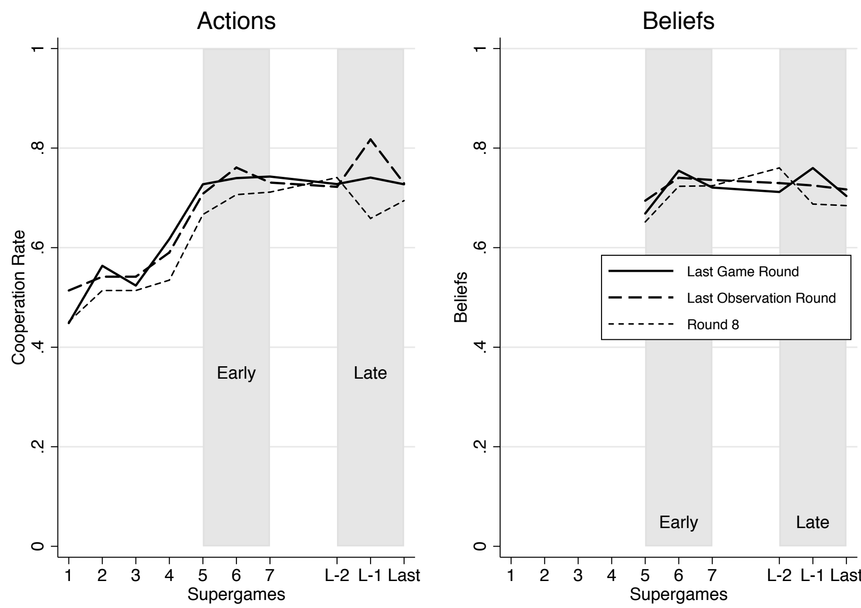

Figure 13: Indefinite Games

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	0.0642	0.187
	(0.294)	(0.281)
Other Cooperated in Previous Supergame	0.251***	0.623***
	(0.0639)	(0.181)
Supergame	0.0133	0.187***
	(0.0430)	(0.0536)
Length of Previous Supergame		-0.00118
		(0.00809)
Cooperated in Supergame 1	2.654***	2.828***
	(0.817)	(0.649)
Constant	0.258	-1.231***
	(0.565)	(0.327)
Observations	1778	1126

Standard errors clustered (at the session level) in parentheses. ***1%, **5%, *10% significance.
All variables refer to behavior in Round 1.

B.3 Beliefs–Comparative Statics

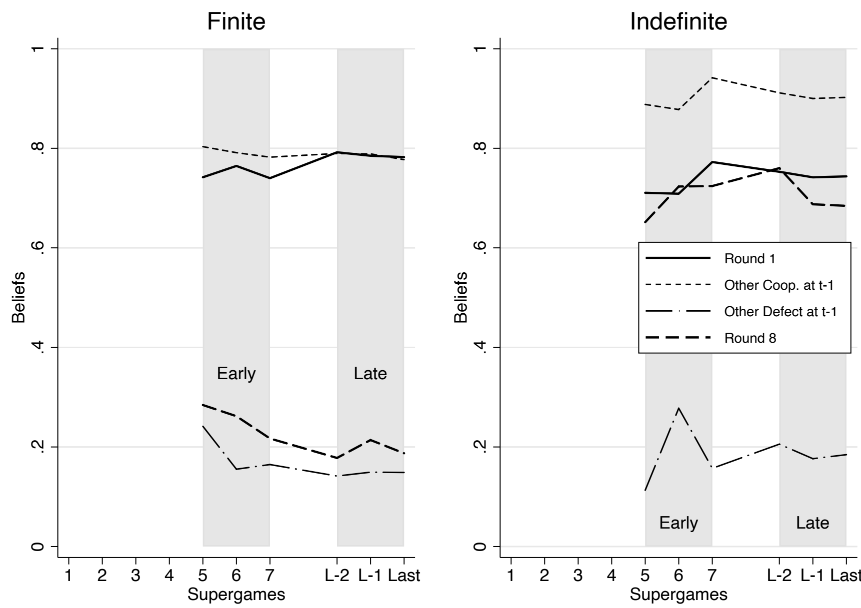

Figure 14: Beliefs Over Supergames

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: $\bar{\mu}(a_j^t|a_j^{t-1} = C)_i^t - \bar{\mu}(a_j^t|a_j^{t-1} = D)_i^t > 0$. Beliefs $\bar{\mu}^T$ in the last round are low in Finite games, but are high in Indefinite games.

B.4 Actions and Beliefs in Round One and Seven

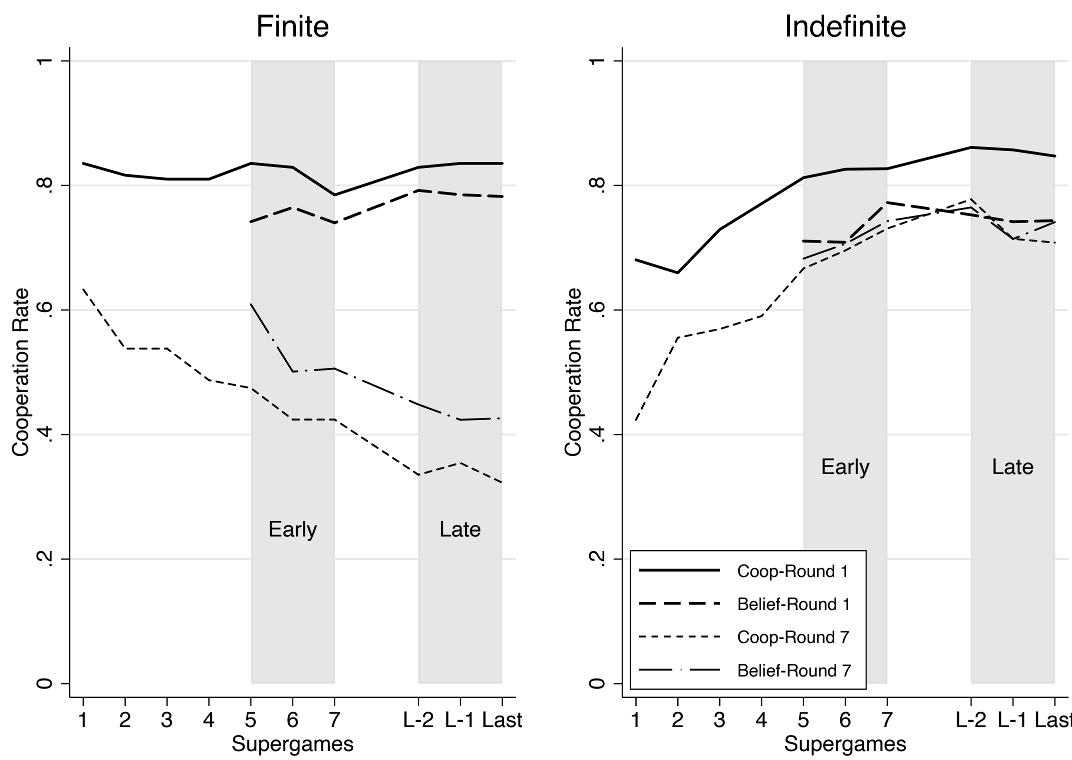

Figure 15: Average Cooperation and Belief

B.5 How Beliefs Relate to Experiences

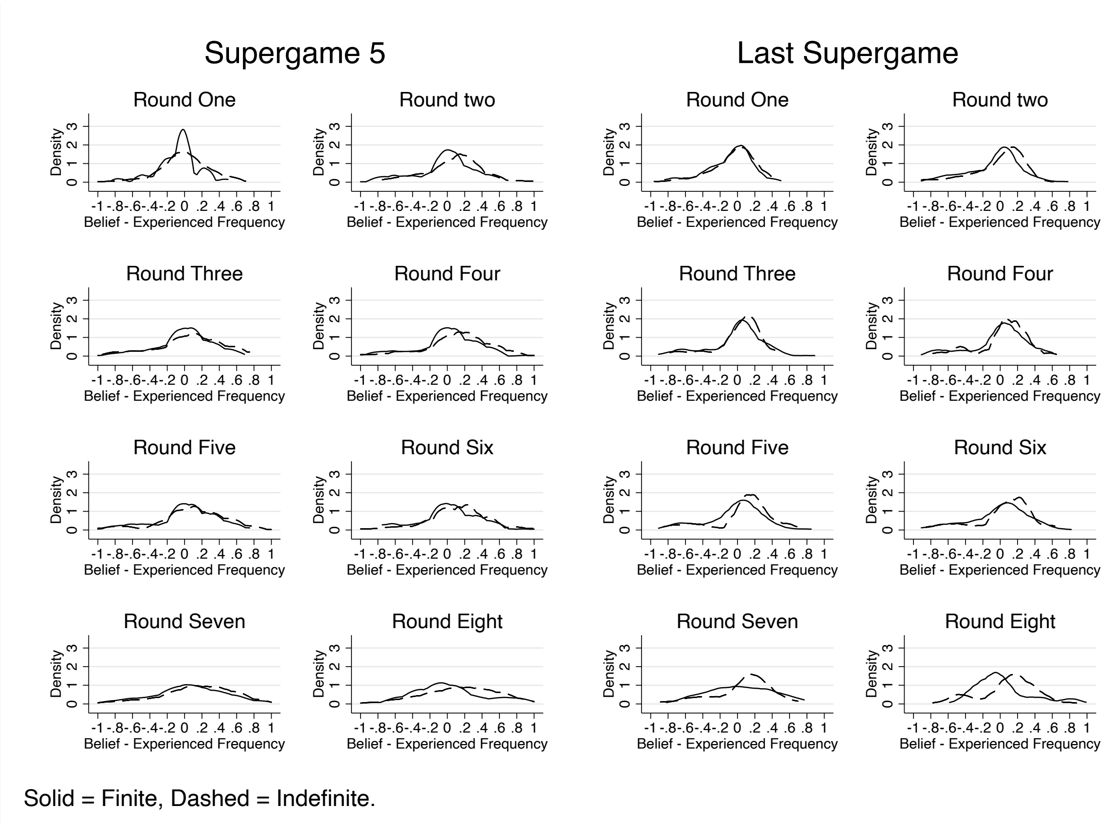

Figure 16: Difference Between Stated Beliefs and Experienced Frequencies of Cooperation by Subject

B.6 Cooperative Path For Early Supergames

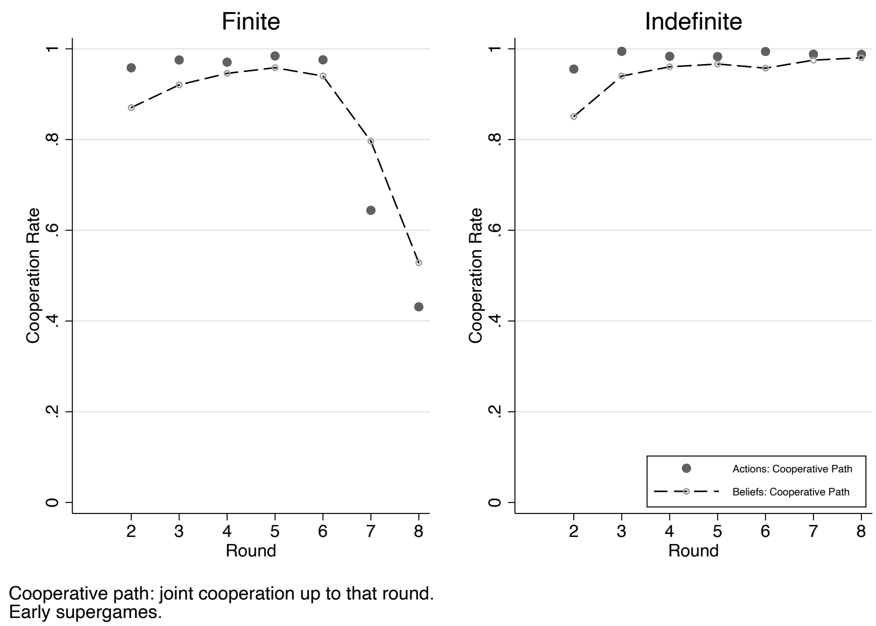

Figure 17: Cooperative Path (First Eight Rounds)

B.7 Accuracy Across Games and Experience

Table 7: 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 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 8: Correlated Random Effects Probit (Marginal Effects)
Dependent Variable: Cooperation

	Finite	Indefinite
Belief	0.462***	0.395***
	(0.0179)	(0.0141)
Round	-0.0336***	-0.00239
	(0.00338)	(0.00281)
Coop. in Round 1, Supergames 1-4	0.244***	0.0808***
	(0.0451)	(0.0248)
Coop. in Last Round, Supergames 1-4	0.126***	0.110***
	(0.0162)	(0.0323)
Observations	3792	3628

Standard errors clustered (at the session level) in parentheses. ***1%, **5%, *10% significance.
Late supergames.

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

Table 9: 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 Complete Estimation Results

Table 10: Finite Games

	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	0.23	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	0.76	0.00	0.06	1.00
AC	0.02	0.01	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
GRIM	0.07	0.02	0.17	0.07	0.00	0.26	0.15	0.23	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.05	*[0.00]*	0.05	0.07	1.00
TFT	0.09	0.12	0.11	0.00	0.00	0.55	0.04	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]*	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	0.14	1.00
T8	0.22	0.20	0.07	0.03	0.00	0.01	0.00	0.51	0.00	0.00	0.00	0.00	0.00	0.00	0.17	0.17	0.01	0.02	0.04	1.00
T7	0.30	0.35	0.00	0.00	0.14	0.00	0.00	0.43	0.43	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	0.01	0.00	*[0.00]*	0.00	0.00	0.48	0.50	0.01	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	0.21	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.07	1.00
ALL			0.06	0.01	0.09	0.11	0.01	0.33	0.20	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 11: Indefinite Games

	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	0.52	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.48	0.00	0.00	0.05	1.00
GRIM	0.15	0.07	0.00	0.09	0.74	0.03	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.06	0.08	0.00	0.00	0.06	1.00
TFT	0.34	0.58	0.07	0.08	0.21	0.27	0.04	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.16	0.14	0.00	0.02	0.01	1.00
STFT	0.04	0.04	0.48	0.00	0.22	0.08	0.01	0.01	0.00	0.00	0.00	0.00	0.03	0.00	0.09	0.00	0.01	0.07	0.08	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	0.00	0.09	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.91	0.00	0.00	0.01	1.00
2TFT	0.05	0.01	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
GRIM3	0.06	0.02	0.00	0.18	0.11	0.19	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.17	0.18	*[0.00]*	0.17	0.01	1.00
ALL			0.16	0.12	0.20	0.18	0.03	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.12	0.17	0.01	0.01		

Estimation on late supergames. SFEM estimate for β is 0.94.

Table 12: 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).

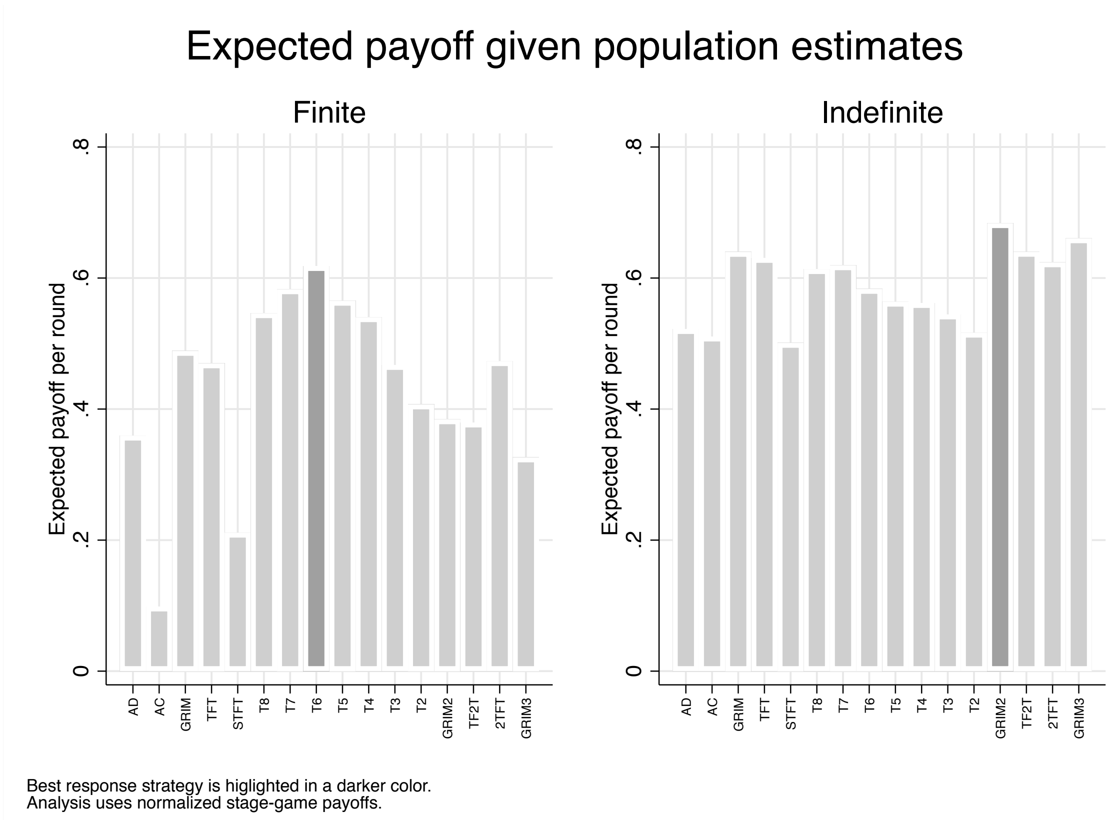

Figure 18: Normalized Expected Payoff by Type Given Strategy Distribution

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.